



US012177532B2

(12) **United States Patent**  
**Ray et al.**

(10) **Patent No.:** **US 12,177,532 B2**  
(45) **Date of Patent:** **Dec. 24, 2024**

(54) **YIELD OPTIMIZATION OF CROSS-SCREEN ADVERTISING PLACEMENT**

(71) Applicant: **VIDEOAMP, INC.**, Los Angeles, CA (US)

(72) Inventors: **Debajyoti Ray**, Marina del Rey, CA (US); **Ross McCray**, Pacific Palisades, CA (US); **David Gullo**, Laguna Hills, CA (US); **Jay Prasad**, San Francisco, CA (US)

(73) Assignee: **VIDEOAMP, INC.**, Los Angeles, CA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **17/075,639**

(22) Filed: **Oct. 20, 2020**

(65) **Prior Publication Data**

US 2021/0105541 A1 Apr. 8, 2021

**Related U.S. Application Data**

(63) Continuation of application No. 15/406,647, filed on Jan. 13, 2017, now Pat. No. 10,812,870, which is a (Continued)

(51) **Int. Cl.**  
**H04N 21/81** (2011.01)  
**G06Q 30/018** (2023.01)  
(Continued)

(52) **U.S. Cl.**  
CPC ..... **H04N 21/812** (2013.01); **G06Q 30/018** (2013.01); **G06Q 30/0244** (2013.01);  
(Continued)

(58) **Field of Classification Search**  
CPC ..... H04N 21/4532; H04N 21/2385; H04N 21/8456; H04N 21/4667; H04N 21/812;  
(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,180,891 A 4/1965 Ellzey, Jr. et al.  
6,009,407 A 12/1999 Garg  
(Continued)

FOREIGN PATENT DOCUMENTS

CN 102077182 A 5/2011  
CN 104054055 A 9/2014  
(Continued)

OTHER PUBLICATIONS

“Adapt.tv Debuts Programmatic TV”, Mar. 6, 2014, Wireless News, 5 Pages, Retrieved from Proquest Technology Collection, <https://search.proquest.com/docview/1504287899?accountid=142944> (retrieved on Sep. 22, 2016).

(Continued)

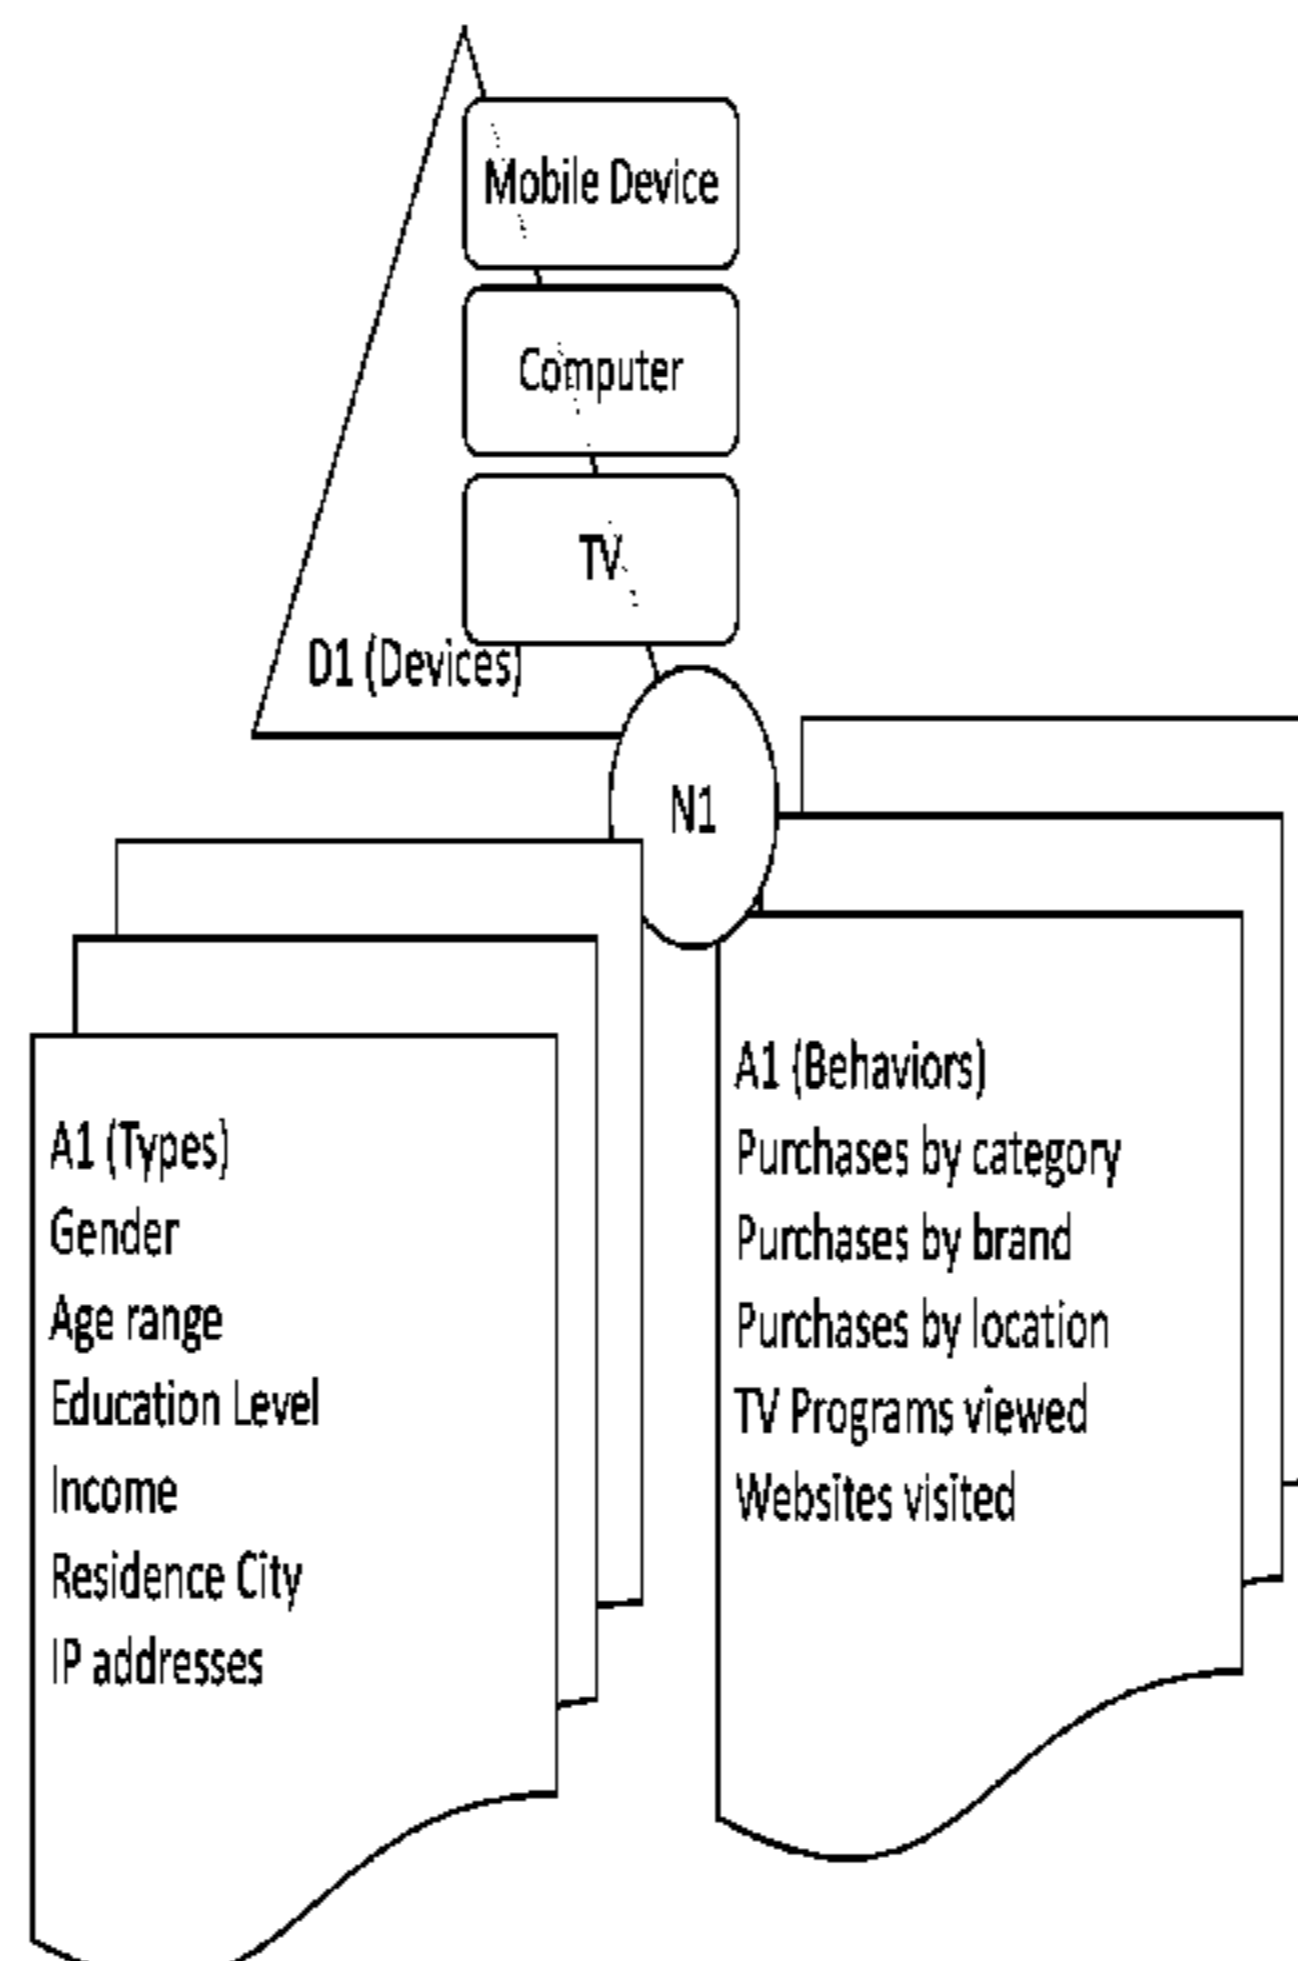
*Primary Examiner* — Adil Ocak

(74) *Attorney, Agent, or Firm* — Buchalter; Jason W. Croft

(57) **ABSTRACT**

The current invention relates to a computer-generated method for optimizing placement of advertising content across multiple different devices. The system permits targeting of advertising content to consumers on TV and mobile devices that can be operated from within a multi-channel video programming distributor’s environment. The system is able to use hard and soft constraints to come up with a number of possible targets for an advertising campaign and can then provide tools for optimizing those targets. The system can allocate advertising campaigns and plans to various inventory types based on the probability of accurate consumer matching. Consumer matching can be achieved by generation of look-alike models in a consumer’s device graph to predict future consumption behavior. The system includes an interface through which a user can adjust various constraints and optimize a distributor’s revenue yield from advertising.

**20 Claims, 13 Drawing Sheets**



**Related U.S. Application Data**

	continuation-in-part of application No. 15/219,262, filed on Jul. 25, 2016, now Pat. No. 9,980,010.	7,661,118 B2	2/2010	Matz et al.	
		7,689,590 B2	3/2010	Fox	
		7,729,940 B2	6/2010	Harvey et al.	
		7,761,346 B2	7/2010	Cooper et al.	
		7,801,778 B2	9/2010	Fox	
		7,802,280 B2	9/2010	Zigmond et al.	
		7,831,685 B2	11/2010	Dai et al.	
		7,835,937 B1	11/2010	Karlsson et al.	
		7,835,939 B1	11/2010	Karlsson	
		7,844,492 B2	11/2010	Perkowski et al.	
		7,870,023 B2	1/2011	Ozer et al.	
		7,904,327 B2	3/2011	Phelan et al.	
		7,917,388 B2	3/2011	van der Riet	
		7,925,549 B2	4/2011	Looney et al.	
		7,945,474 B1	5/2011	Koinov	
		7,945,928 B2	5/2011	Matz et al.	
		7,949,561 B2	5/2011	Briggs	
		7,949,562 B2	5/2011	Collins	
		7,949,563 B2	5/2011	Collins	
		7,974,889 B2	7/2011	Raimbeault	
		8,000,993 B2	8/2011	Harvey et al.	
		8,006,261 B1	8/2011	Haberman et al.	
		8,046,797 B2 *	10/2011	Bentolila .....	H04N 21/4663 725/46
		8,060,398 B2	11/2011	Canning et al.	
		8,069,142 B2	11/2011	Davis et al.	
		8,090,613 B2	1/2012	Kalb et al.	
		8,099,316 B2	1/2012	Moukas et al.	
		8,099,317 B2	1/2012	Moukas et al.	
		8,099,318 B2	1/2012	Moukas et al.	
		8,099,323 B2	1/2012	Veeraraghavan	
		8,112,301 B2	2/2012	Harvey et al.	
		8,117,050 B2	2/2012	Li et al.	
		8,117,067 B2 *	2/2012	Ketchum .....	G06Q 40/00 705/14.4
		8,135,626 B2	3/2012	Das et al.	
		8,160,916 B2	4/2012	Moukas et al.	
		8,175,914 B1	5/2012	Benson et al.	
		8,180,891 B1	5/2012	Harrison	
		8,195,508 B1	6/2012	Calder et al.	
		8,239,242 B2	8/2012	Moukas et al.	
		8,249,925 B2	8/2012	Broms et al.	
		8,255,949 B1	8/2012	Bayer et al.	
		8,265,996 B2	9/2012	Steelberg et al.	
		8,296,793 B2	10/2012	Johnson	
		8,321,273 B2	11/2012	Briggs	
		8,321,274 B2	11/2012	Collins et al.	
		8,321,275 B2	11/2012	Collins et al.	
		8,341,047 B1	12/2012	Furney et al.	
		8,346,593 B2	1/2013	Fanelli et al.	
		8,346,607 B1	1/2013	Benson et al.	
		8,352,307 B2	1/2013	Moukas et al.	
		8,352,320 B2	1/2013	MacDonald et al.	
		8,386,310 B2	2/2013	Weyer et al.	
		8,386,398 B1	2/2013	Hickman	
		8,392,252 B2	3/2013	Kaufman et al.	
		8,401,887 B2	3/2013	Moukas et al.	
		8,402,482 B2	3/2013	Woodward et al.	
		8,417,556 B2	4/2013	Moukas et al.	
		8,423,408 B1	4/2013	Barnes et al.	
		8,442,859 B1	5/2013	Broms et al.	
		8,452,797 B1	5/2013	Paleja et al.	
		8,473,339 B1	6/2013	McKennon et al.	
		8,478,675 B1	7/2013	Walia et al.	
		8,489,460 B2	7/2013	Kamath	
		8,495,679 B2 *	7/2013	Labeeb .....	H04N 7/163 725/46
		8,504,416 B2	8/2013	Calder et al.	
		8,516,515 B2	8/2013	Zigmond et al.	
		8,527,341 B2	9/2013	Feuerstein et al.	
		8,527,342 B2	9/2013	Feuerstein et al.	
		8,533,048 B2 *	9/2013	Charania .....	G06Q 30/0247 705/14.69
		8,533,049 B2	9/2013	Flake et al.	
		8,539,067 B2	9/2013	Juda et al.	
		8,539,072 B1	9/2013	Harrison	
		8,554,619 B2	10/2013	Juda et al.	
		8,560,385 B2	10/2013	Atazky et al.	
		8,566,224 B2	10/2013	Walia et al.	

(60) Provisional application No. 62/290,387, filed on Feb. 2, 2016, provisional application No. 62/278,888, filed on Jan. 14, 2016, provisional application No. 62/264,764, filed on Dec. 8, 2015, provisional application No. 62/196,592, filed on Jul. 24, 2015.

- (51) **Int. Cl.**  
**G06Q 30/0242** (2023.01)  
**G06Q 30/0251** (2023.01)  
**G06Q 30/0273** (2023.01)  
**H04N 21/2385** (2011.01)  
**H04N 21/25** (2011.01)  
**H04N 21/2543** (2011.01)  
**H04N 21/258** (2011.01)  
**H04N 21/266** (2011.01)  
**H04N 21/2668** (2011.01)  
**H04N 21/45** (2011.01)  
**H04N 21/466** (2011.01)  
**H04N 21/845** (2011.01)

(52) **U.S. Cl.**  
 CPC ..... **G06Q 30/0264** (2013.01); **G06Q 30/0267** (2013.01); **G06Q 30/0269** (2013.01); **G06Q 30/0275** (2013.01); **H04N 21/2385** (2013.01); **H04N 21/252** (2013.01); **H04N 21/2543** (2013.01); **H04N 21/25883** (2013.01); **H04N 21/25891** (2013.01); **H04N 21/266** (2013.01); **H04N 21/2668** (2013.01); **H04N 21/4532** (2013.01); **H04N 21/4667** (2013.01); **H04N 21/8456** (2013.01)

(58) **Field of Classification Search**  
 CPC ..... H04N 21/252; H04N 21/2543; H04N 21/25883; H04N 21/25891; H04N 21/266; H04N 21/2668; G06Q 30/018; G06Q 30/0264; G06Q 30/0267; G06Q 30/0269; G06Q 30/0244; G06Q 30/0275  
 See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,055,573 A	4/2000	Gardenswartz et al.
6,061,691 A	5/2000	Fox
6,119,101 A	9/2000	Peckover
6,253,187 B1	6/2001	Fox
6,286,005 B1	9/2001	Cannon
6,298,330 B1	10/2001	Gardenswartz et al.
6,567,824 B2	5/2003	Fox
6,868,389 B1	3/2005	Wilkins et al.
6,895,387 B1	5/2005	Roberts et al.
6,920,464 B2	7/2005	Fox
6,937,996 B1	8/2005	Forsythe et al.
6,993,494 B1	1/2006	Boushy et al.
7,062,510 B1	6/2006	Eldering
7,124,091 B1	10/2006	Khoo et al.
7,136,871 B2	11/2006	Ozer et al.
7,158,943 B2	1/2007	van der Riet
7,222,099 B2	5/2007	Forsythe et al.
7,228,287 B1	6/2007	Samson et al.
7,356,547 B2	4/2008	Ozer et al.
7,363,254 B2	4/2008	Skinner
7,423,406 B2	9/2008	Geniusz
7,444,658 B1	10/2008	Matz et al.
7,536,316 B2	5/2009	Ozer et al.
7,548,929 B2	6/2009	Collins et al.
7,574,404 B2	8/2009	Forsythe et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

8,566,861 B2	10/2013	Zigmond et al.	2007/0169143 A1	7/2007	Li
8,583,659 B1	11/2013	Alexandrescu et al.	2007/0198355 A1	8/2007	Samson et al.
8,589,233 B2	11/2013	Flake et al.	2008/0167943 A1	7/2008	O'Neil et al.
8,620,795 B2	12/2013	Veeraraghavan	2008/0221949 A1*	9/2008	Delurgio ..... G06Q 10/063 705/7.11
8,630,902 B2	1/2014	Baum et al.	2008/0271070 A1*	10/2008	Kanojia ..... H04N 21/25841 725/32
8,639,567 B2	1/2014	Winters	2009/0037252 A1	2/2009	Joyce
8,645,186 B2	2/2014	Moukas et al.	2009/0091571 A1	4/2009	Zalewski
8,650,066 B2	2/2014	Shah et al.	2009/0150215 A1	6/2009	Kalb et al.
8,655,695 B1	2/2014	Qu et al.	2009/0150224 A1	6/2009	Lu et al.
8,655,730 B1	2/2014	Swan et al.	2009/0171780 A1	7/2009	Aldrey et al.
8,671,016 B2	3/2014	Little	2009/0172728 A1	7/2009	Shkedi et al.
8,700,452 B1	4/2014	McKennon et al.	2009/0199107 A1*	8/2009	Lewis ..... G06Q 30/0241 715/745
8,700,462 B2	4/2014	Collins	2009/0248478 A1	10/2009	Duggal et al.
8,732,015 B1	5/2014	Beckerman et al.	2009/0248680 A1	10/2009	Kalavade
8,738,433 B2	5/2014	DiOrio et al.	2009/0298480 A1*	12/2009	Khambete ..... G06Q 30/0254 455/414.1
8,738,446 B1	5/2014	L'Heureux et al.	2010/0017298 A1	1/2010	Stukenborg et al.
8,744,908 B2	6/2014	Kalb et al.	2010/0057546 A1	3/2010	Wang et al.
8,751,359 B2	6/2014	Walia et al.	2010/0057560 A1*	3/2010	Skudlark ..... H04N 21/812 705/14.49
8,751,418 B1	6/2014	Sutter et al.	2010/0121671 A1	5/2010	Boutilier et al.
8,752,086 B2*	6/2014	Conant ..... G06Q 30/0272 725/34	2010/0191600 A1*	7/2010	Sideman ..... G06Q 30/02 705/14.53
8,776,111 B1*	7/2014	Eldering ..... H04N 21/25891 725/34	2010/0235219 A1	9/2010	Merrick et al.
8,782,683 B2*	7/2014	Balakrishnan ..... H04N 21/433 725/14	2010/0250332 A1	9/2010	Ghosh et al.
8,788,345 B2	7/2014	Kamath	2010/0324982 A1	12/2010	Marolli
8,819,249 B2	8/2014	Harrison	2010/0324992 A1*	12/2010	Birch ..... G06Q 30/02 705/14.49
8,819,255 B1	8/2014	Harrison	2011/0035277 A1	2/2011	Kodialam et al.
8,825,677 B2	9/2014	Gausebeck et al.	2011/0040611 A1	2/2011	Simmons et al.
8,831,987 B2	9/2014	Knapp et al.	2011/0040612 A1	2/2011	Simmons et al.
8,904,021 B2	12/2014	Harrison	2011/0040613 A1	2/2011	Simmons et al.
8,909,771 B2*	12/2014	Heath ..... G06Q 30/02 709/224	2011/0040635 A1	2/2011	Simmons et al.
8,924,465 B1	12/2014	Tunguz-Zawislak	2011/0040636 A1	2/2011	Simmons et al.
8,935,721 B2*	1/2015	Tidwell ..... G06Q 30/0244 725/34	2011/0099048 A1	4/2011	Weiss et al.
8,949,890 B2	2/2015	Evans et al.	2011/0106618 A1	5/2011	Ben-Moshe
8,972,281 B2	3/2015	Pokonosky	2011/0119126 A1	5/2011	Park et al.
8,983,859 B2	3/2015	Nice et al.	2011/0119128 A1	5/2011	Fang et al.
8,990,105 B1	3/2015	Shatkin-Margolis et al.	2011/0225037 A1	9/2011	Tunca et al.
9,026,886 B1	5/2015	Yeo	2011/0225604 A1	9/2011	Bova
9,053,497 B2	6/2015	Benyamin	2011/0231242 A1	9/2011	Dilling et al.
9,154,942 B2	10/2015	Harrison et al.	2011/0231243 A1	9/2011	Bhatia et al.
9,167,419 B2	10/2015	Harrison	2011/0231245 A1	9/2011	Bhatia et al.
9,980,010 B2	5/2018	Ray et al.	2011/0231246 A1	9/2011	Bhatia et al.
9,980,011 B2	5/2018	Ray et al.	2011/0231253 A1	9/2011	Crawford et al.
10,085,073 B2	9/2018	Ray et al.	2011/0231253 A1	9/2011	Dilling et al.
10,123,063 B1	11/2018	Brown	2011/0231264 A1	9/2011	Haumont et al.
10,136,174 B2	11/2018	Ray et al.	2011/0238485 A1	9/2011	Haumont et al.
10,417,658 B1	9/2019	Tsemekhman et al.	2011/0239243 A1*	9/2011	Dierks ..... H04N 21/812 725/24
10,555,050 B2	2/2020	Ray et al.	2011/0288907 A1	11/2011	Harvey et al.
10,667,020 B2	5/2020	Ray et al.	2012/0005216 A1	1/2012	Moritz et al.
10,812,870 B2	10/2020	Ray et al.	2012/0022937 A1	1/2012	Bhatia et al.
11,425,441 B2	8/2022	Ray et al.	2012/0059713 A1	3/2012	Galas et al.
11,463,786 B2	10/2022	Ray et al.	2012/0089455 A1	4/2012	Belani et al.
11,856,272 B2	12/2023	Ray et al.	2012/0095863 A1	4/2012	Schiff et al.
2002/0072966 A1	6/2002	Eldering et al.	2012/0120250 A1	5/2012	Shintani et al.
2002/0151992 A1	10/2002	Hoffberg et al.	2012/0158485 A1	6/2012	Ogawa
2002/0178447 A1	11/2002	Plotnick et al.	2012/0158486 A1	6/2012	Ogawa
2002/0184047 A1	12/2002	Plotnick et al.	2012/0158487 A1	6/2012	Ogawa
2003/0066068 A1*	4/2003	Gutta ..... H04N 21/4661 725/9	2012/0191529 A1	7/2012	Lewis et al.
2004/0098744 A1*	5/2004	Gutta ..... H04N 21/47 725/46	2012/0239809 A1	9/2012	Mazumdar et al.
2004/0133909 A1	7/2004	Ma	2012/0253928 A1	10/2012	Jackson et al.
2005/0021397 A1	1/2005	Cui et al.	2012/0271708 A1	10/2012	Ogawa
2005/0137958 A1	6/2005	Huber et al.	2012/0271709 A1	10/2012	Ogawa
2006/0190330 A1	8/2006	Tollinger et al.	2012/0284212 A1	11/2012	Lin et al.
2006/0217110 A1	9/2006	Othmer	2012/0303429 A1*	11/2012	Nolledo ..... G06Q 30/02 705/14.16
2006/0253328 A1	11/2006	Kohli et al.	2012/0310745 A1	12/2012	Bhatia et al.
2007/0028266 A1*	2/2007	Trajkovic ..... H04N 21/6582 725/46	2012/0323674 A1	12/2012	Simmons et al.
2007/0038509 A1	2/2007	Jain et al.	2013/0006754 A1	1/2013	Horvitz et al.
2007/0143186 A1	6/2007	Apple et al.	2013/0054349 A1	2/2013	Ogawa
			2013/0124309 A1*	5/2013	Traasdahl ..... G06Q 30/0241 705/14.49
			2013/0185750 A1	7/2013	Ayoub
			2013/0205339 A1	8/2013	Haberman et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

2013/0262225 A1 10/2013 Ahrens et al.  
 2013/0346182 A1 12/2013 Cheng et al.  
 2014/0025660 A1 1/2014 Mohammed et al.  
 2014/0040008 A1 2/2014 Belani et al.  
 2014/0046777 A1 2/2014 Markey et al.  
 2014/0052527 A1 2/2014 Roundtree  
 2014/0058826 A1 2/2014 Ogawa  
 2014/0095320 A1 4/2014 Sivaramakrishnan et al.  
 2014/0099623 A1 4/2014 Amit et al.  
 2014/0109123 A1 4/2014 Balakrishnan et al.  
 2014/0114773 A1 4/2014 Stukenborg et al.  
 2014/0136338 A1 5/2014 Ringdahl  
 2014/0157309 A1\* 6/2014 Small ..... H04H 60/66  
 725/35  
 2014/0278912 A1 9/2014 Hughes et al.  
 2014/0278921 A1 9/2014 Sankaran et al.  
 2014/0278981 A1 9/2014 Mersov et al.  
 2014/0279036 A1 9/2014 Ke et al.  
 2014/0289017 A1 9/2014 Trenkle et al.  
 2014/0289765 A1 9/2014 Kitts et al.  
 2014/0297400 A1 10/2014 Sandholm  
 2015/0025936 A1 1/2015 Garel et al.  
 2015/0066662 A1 3/2015 Knapp et al.  
 2015/0088635 A1 3/2015 Maycotte et al.  
 2015/0112803 A1 4/2015 Heiser et al.  
 2015/0121418 A1\* 4/2015 Jain ..... H04N 21/812  
 725/32  
 2015/0161186 A1 6/2015 Vemuri  
 2015/0186945 A1\* 7/2015 Samuel ..... G06Q 30/0267  
 705/14.25  
 2015/0208110 A1 7/2015 Small et al.  
 2015/0235275 A1 8/2015 Shah et al.  
 2015/0287097 A1 10/2015 Umeda  
 2015/0339702 A1 11/2015 Lin et al.  
 2015/0350320 A1 12/2015 Yang  
 2015/0358818 A1\* 12/2015 Dipaola ..... G06Q 30/0267  
 726/4  
 2016/0092933 A1 3/2016 Xu et al.  
 2016/0112735 A1\* 4/2016 Stephens, Jr. .... H04N 21/466  
 725/14  
 2016/0117718 A1 4/2016 Hood et al.  
 2016/0117740 A1\* 4/2016 Linden ..... G06Q 50/01  
 705/14.66  
 2016/0119689 A1\* 4/2016 Hood ..... H04N 21/25891  
 725/14  
 2016/0125461 A1 5/2016 Sivaramakrishnan et al.  
 2016/0267523 A1\* 9/2016 Biswas ..... G06Q 30/0201  
 2017/0024765 A1 1/2017 Barenholz et al.  
 2017/0034591 A1 2/2017 Ray et al.  
 2017/0034592 A1 2/2017 Ray et al.  
 2017/0034593 A1 2/2017 Ray et al.  
 2017/0083941 A1\* 3/2017 Biswas ..... G06Q 30/0275  
 2017/0099525 A1 4/2017 Ray et al.  
 2017/0142465 A1 5/2017 Ray et al.  
 2018/0338189 A1 11/2018 Ray et al.  
 2018/0376220 A1 12/2018 Ray et al.  
 2019/0149869 A1 5/2019 Ray et al.  
 2019/0373297 A1\* 12/2019 Sarkhel ..... G06N 5/003  
 2020/0294086 A1 9/2020 Traasdahl et al.  
 2021/0185408 A1 6/2021 Ray et al.  
 2023/0032230 A1 2/2023 Ray et al.

FOREIGN PATENT DOCUMENTS

WO 2017019643 A1 2/2017  
 WO 2017019645 A1 2/2017  
 WO 2017019646 A1 2/2017  
 WO 2017019647 A1 2/2017  
 WO 2017124041 A1 7/2017  
 WO 2018022144 A1 2/2018

OTHER PUBLICATIONS

Brightroll Inc., “How is Ad Inventory like Fruit?”, 5 pages, Jun. 29, 2016, available at <https://brightroll.com/insights/how-ad-inventory-fruit> (last accessed Jul. 26, 2017).  
 Comscore Inc., “Cross Platform Future in Focus 2016”, Whitepaper, 55 pages, Mar. 2016, available at <https://www.comscore.com/insights/Presentations-and-Whitepapers/2016/2016-US-cross-Platform-Future-in-Focus> (last accessed Jul. 26, 2017).  
 “Cross-Screen Planning Whitepaper”, 13 pages. 2016, available from <https://www.tubemogul.com/> (last accessed Jul. 26, 2017).  
 “Drawbridge Connected Consumer Graph”, The Largest Independent Cross-Device Identity Solution, 7 pages, 2016, available at <https://drawbridge.com/c/graph> (last accessed Jul. 26, 2017).  
 Final Office Action for U.S. Appl. No. 16/782,006, dated Nov. 17, 2022, 20 pages.  
 Final Office Action for U.S. Appl. No. 16/782,006, dated Oct. 25, 2023, 10 pages.  
 Freewheel: “Video Monetization Report”, Feb. 2016, 40 Pages, Available at <https://freewheel.tv/freewheel-views/a-peel-inside-the-freewheel-video-monetization-report-q2-2016/> (last accessed Jul. 26, 2017).  
 International Search Report and Written Opinion for PCT/US2016/043952 mailed Oct. 6, 2016, 8 pages.  
 International Search Report and Written Opinion for PCT/US2016/043954 mailed Dec. 12, 2016, 13 pages.  
 International Search Report and Written Opinion for PCT/US2016/043955 mailed Nov. 29, 2016, 19 pages.  
 International Search Report and Written Opinion for PCT/US2016/043958 mailed Nov. 29, 2016, 12 pages.  
 International Search Report and Written Opinion for PCT/US2017/013569 mailed Apr. 14, 2017, 13 pages.  
 International Search Report and Written Opinion for PCT/US2017/015870 mailed May 19, 2017, 17 pages.  
 Non-Final Office Action for U.S. Appl. No. 15/219,259, dated Jun. 15, 2017, 13 pages.  
 Non-Final Office Action for U.S. Appl. No. 15/219,262, dated Jun. 1, 2017, 30 pages.  
 Non-Final Office Action for U.S. Appl. No. 15/219,264, dated Apr. 19, 2017, 34 pages.  
 Non-Final Office Action for U.S. Appl. No. 15/219,268, dated Jun. 13, 2017, 20 pages.  
 Non-Final Office Action for U.S. Appl. No. 16/782,006, dated Apr. 11, 2023, 19 pages.  
 Notice of Allowance for U.S. Appl. No. 15/985,639, dated Nov. 29, 2023, 10 pages.  
 Notice of Allowance for U.S. Appl. No. 16/190,056, dated Apr. 13, 2022, 8 pages.  
 Notice of Allowance for U.S. Appl. No. 17/938,033, dated Nov. 17, 2023, 8 Pages.  
 Office Action for Chinese Application No. 201680055868.9, dated May 13, 2021, 16 pages.  
 Supplementary European Search Report for European Patent Application No. 16831218.9, dated Jan. 31, 2019, 9 Pages.  
 Supplementary European Search Report for European Patent Application No. 16831220.5, dated Apr. 18, 2019, 15 Pages.  
 Supplementary European Search Report for European Patent Application No. 16831221.3, dated Feb. 12, 2019, 9 Pages.  
 Supplementary European Search Report for European Patent Application No. 16831222.1, dated Feb. 11, 2019, 10 Pages.  
 Supplementary European Search Report in European Patent Application No. 17739108.3, dated Jul. 31, 2019, 10 Pages.  
 TAPAD: “Measuring Cross-Device, The Methodology”, whitepaper, 5 pages, Mar. 2017, available at <https://www.tapad.com/resources/cross-device-the-methodology>, (last accessed Jul. 26, 2017).  
 “The Device Graph”, Tapad, Available at <https://www.tapad.com/device-graph>, last accessed Jul. 26, 2017, 7 Pages.  
 Tubemogul Inc., “Programmatic Advertising Software”, TubeMogul Software, 8 pages, Jul. 15, 2016, available from <https://www.tubemogul.com/tubemogul-software/> (last accessed Jul. 26, 2017).  
 U.S. Appl. No. 16/782,006, filed Feb. 4, 2020, titled “Cross-Screen Measurement Accuracy in Advertising Performance”, 61 Pages.  
 Notice of Allowance for U.S. Appl. No. 15/985,639, dated Apr. 5, 2024, 8 pages.

(56)

**References Cited**

OTHER PUBLICATIONS

Notice of Allowance for U.S. Appl. No. 16/782,006, dated Apr. 11, 2024, 08 pages.

Notice of Allowance for U.S. Appl. No. 17/938,033, dated Mar. 20, 2024, 8 pages.

Office Action of Chinese Patent Application No. 2021110781106 dated Sep. 25, 2024, 8 pgs with translations.

\* cited by examiner

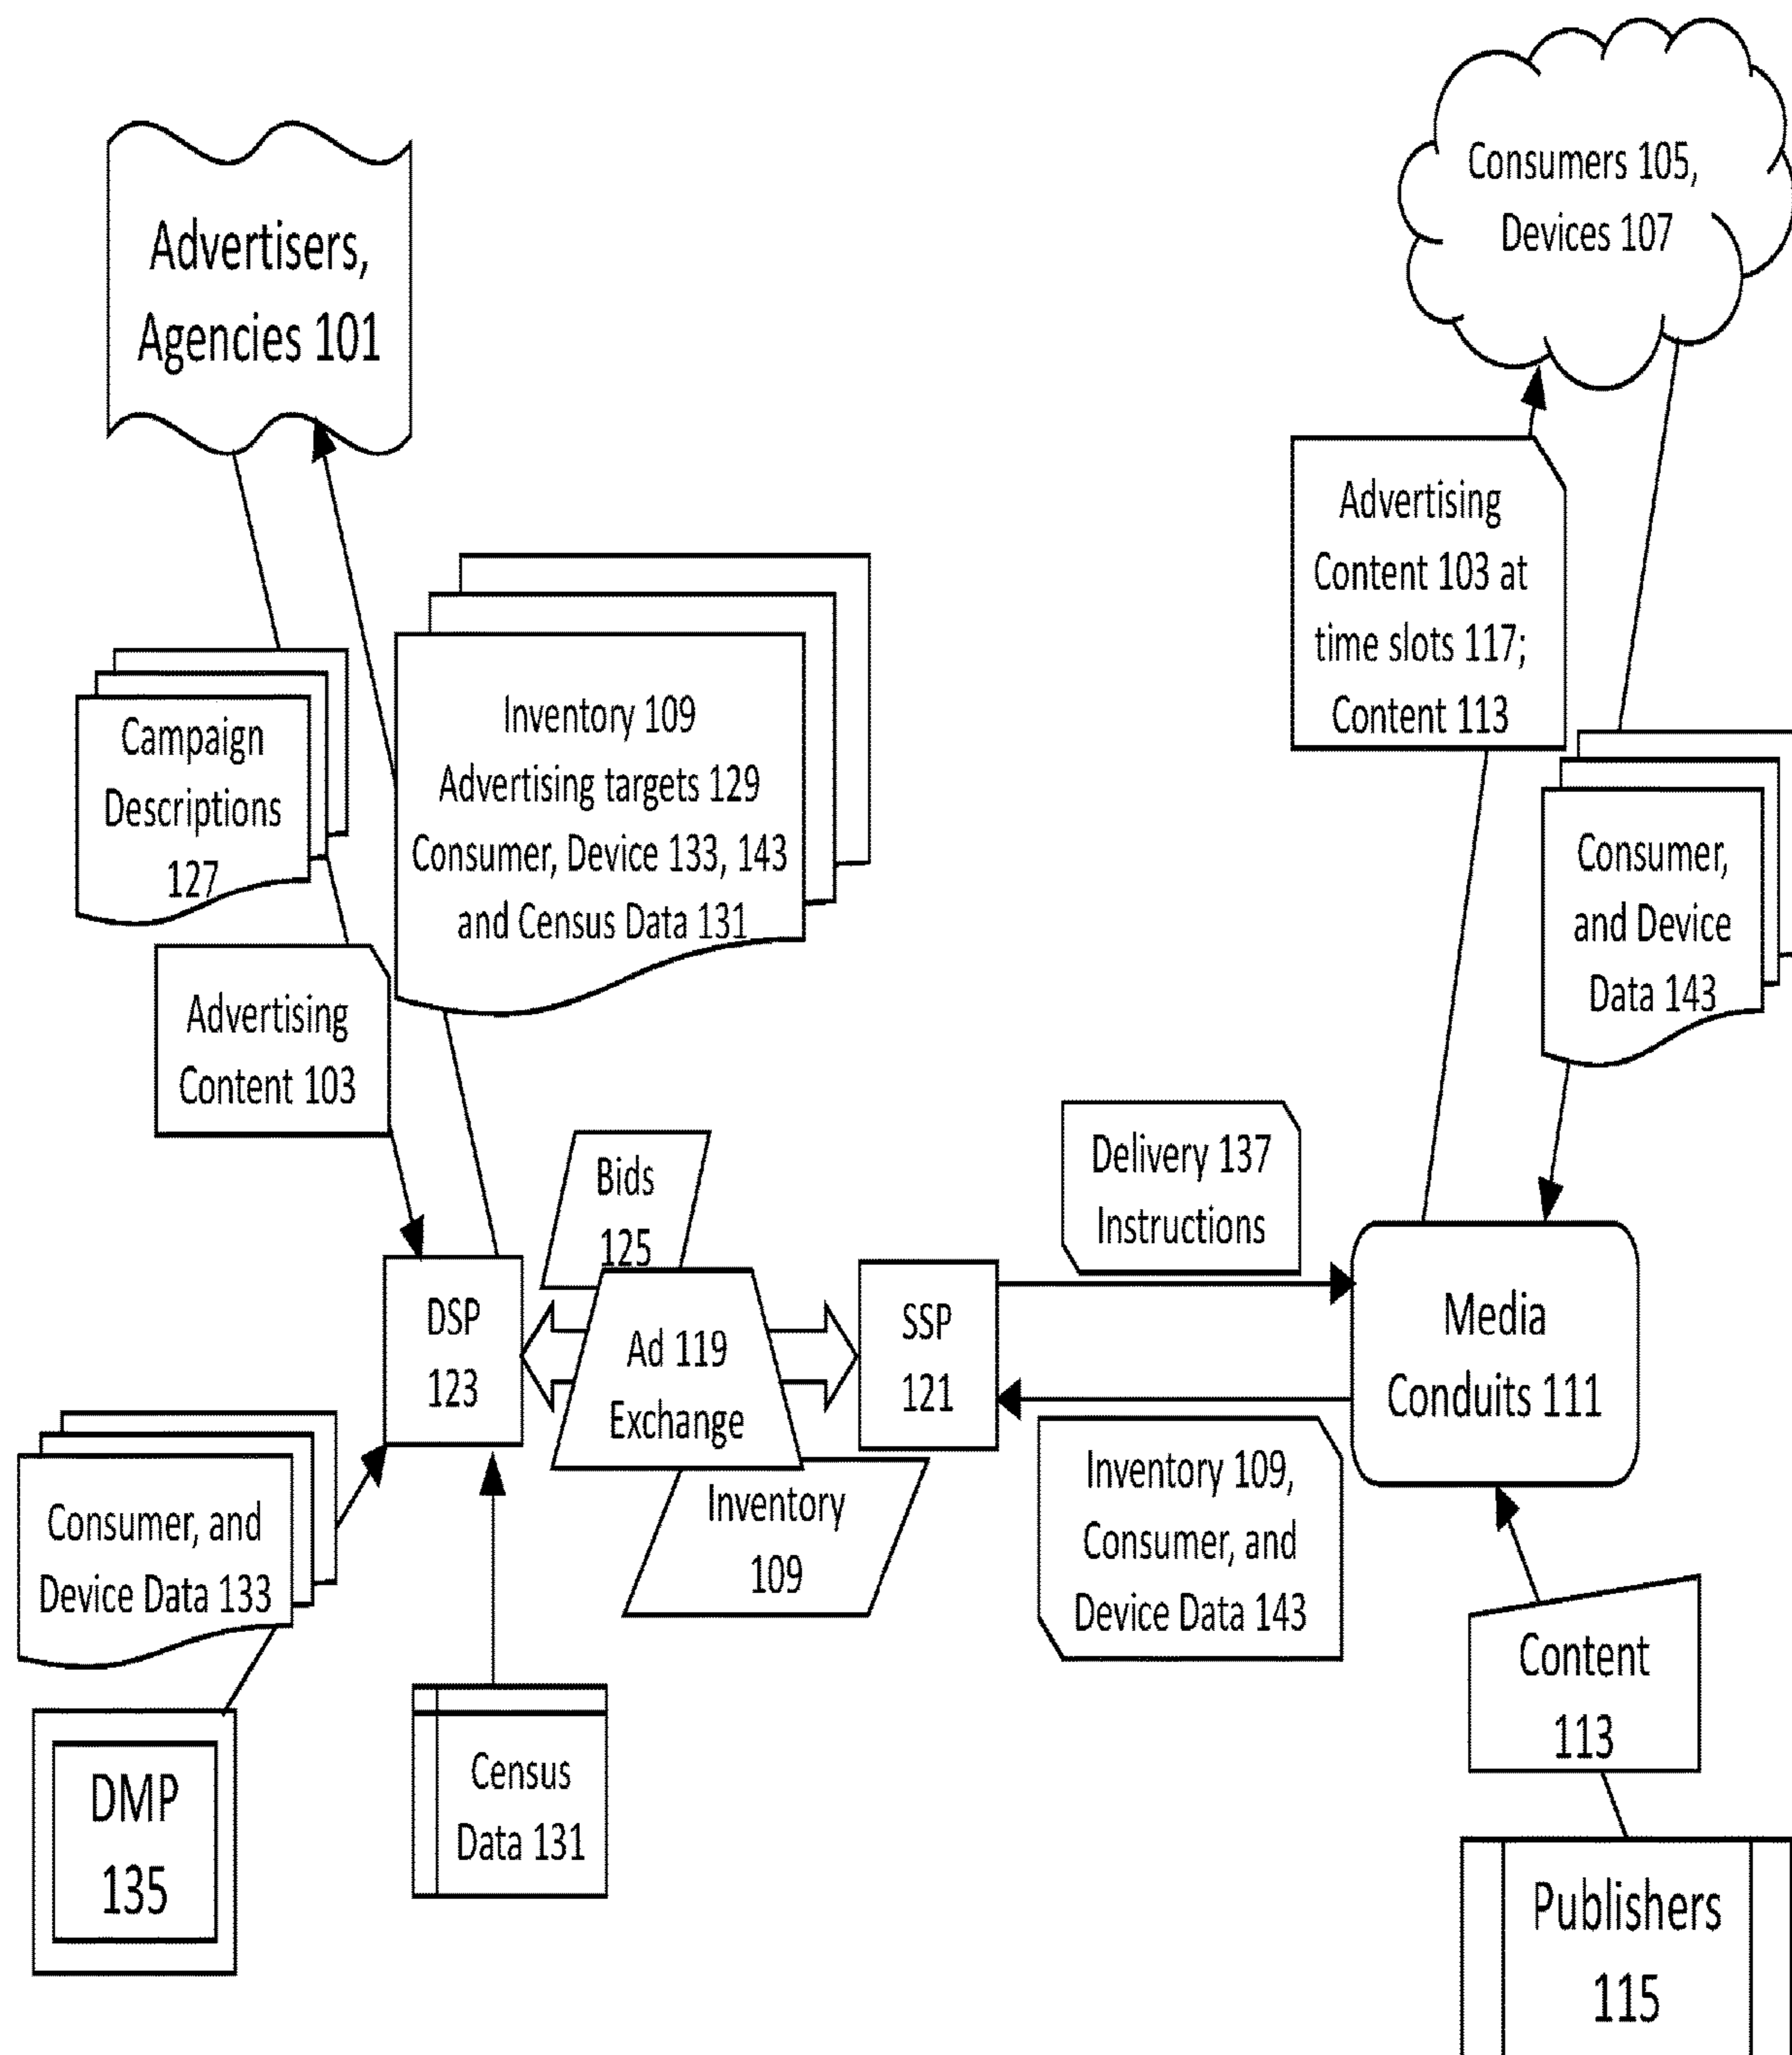
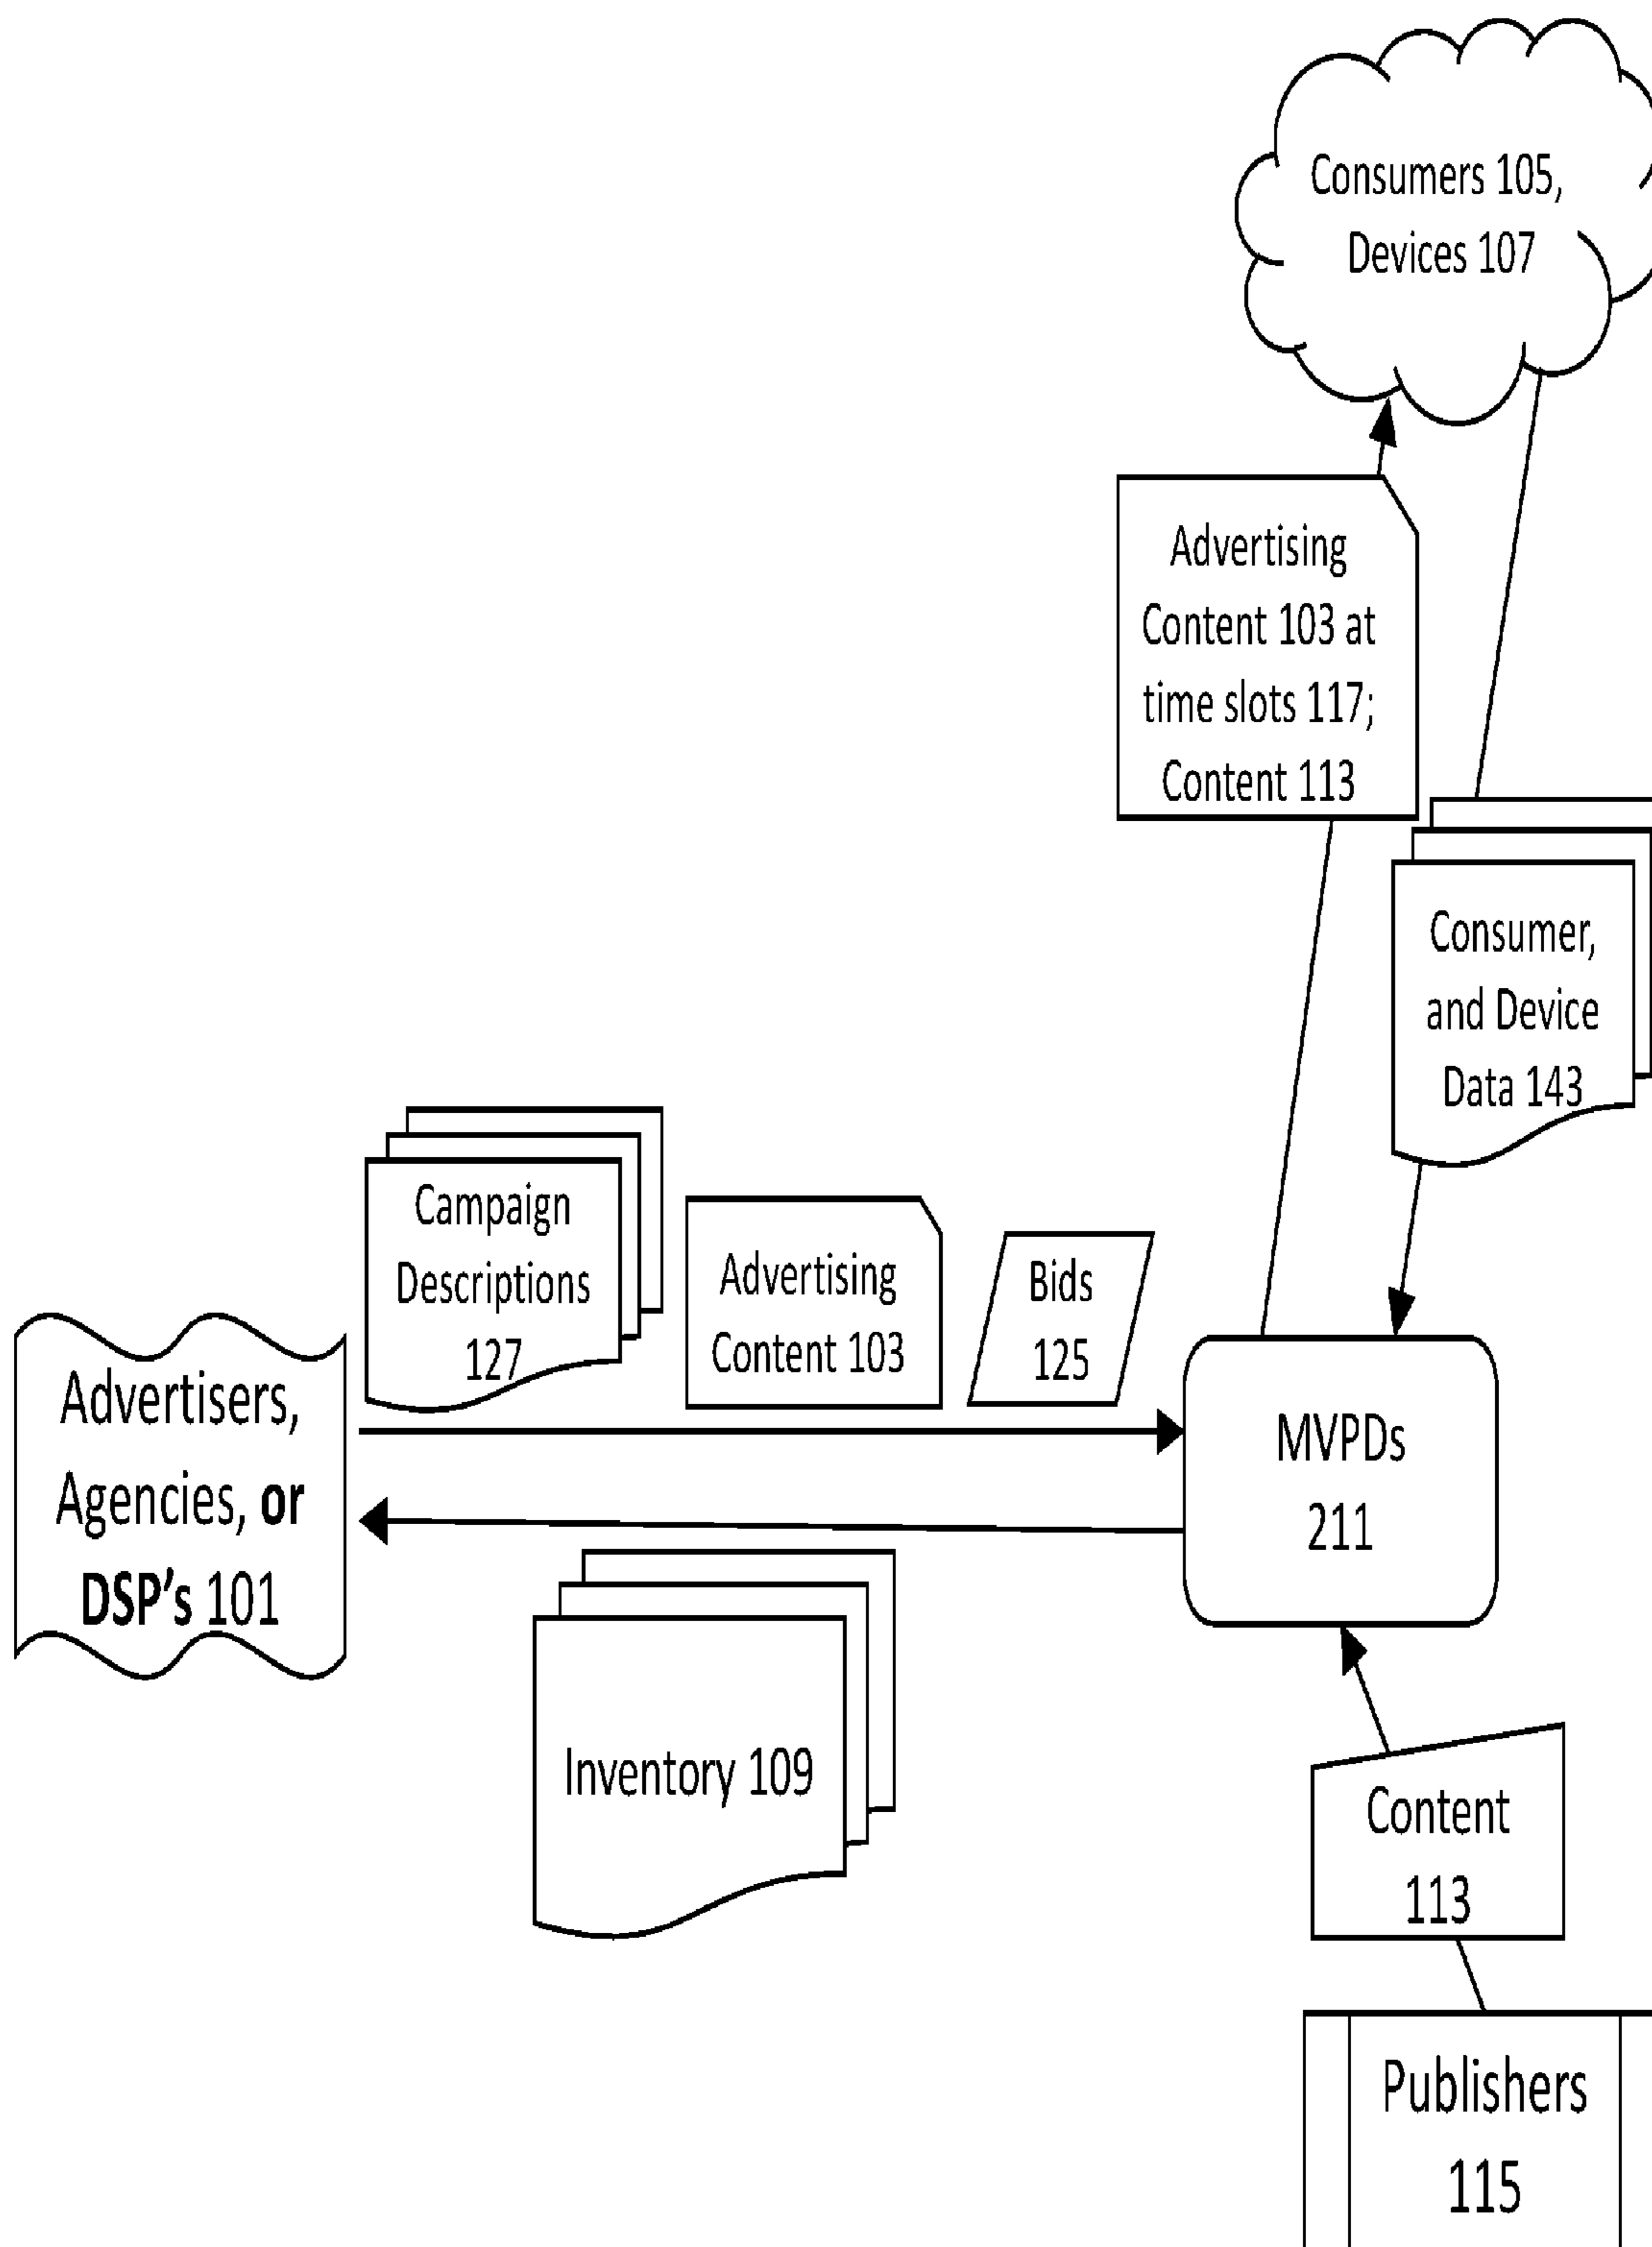
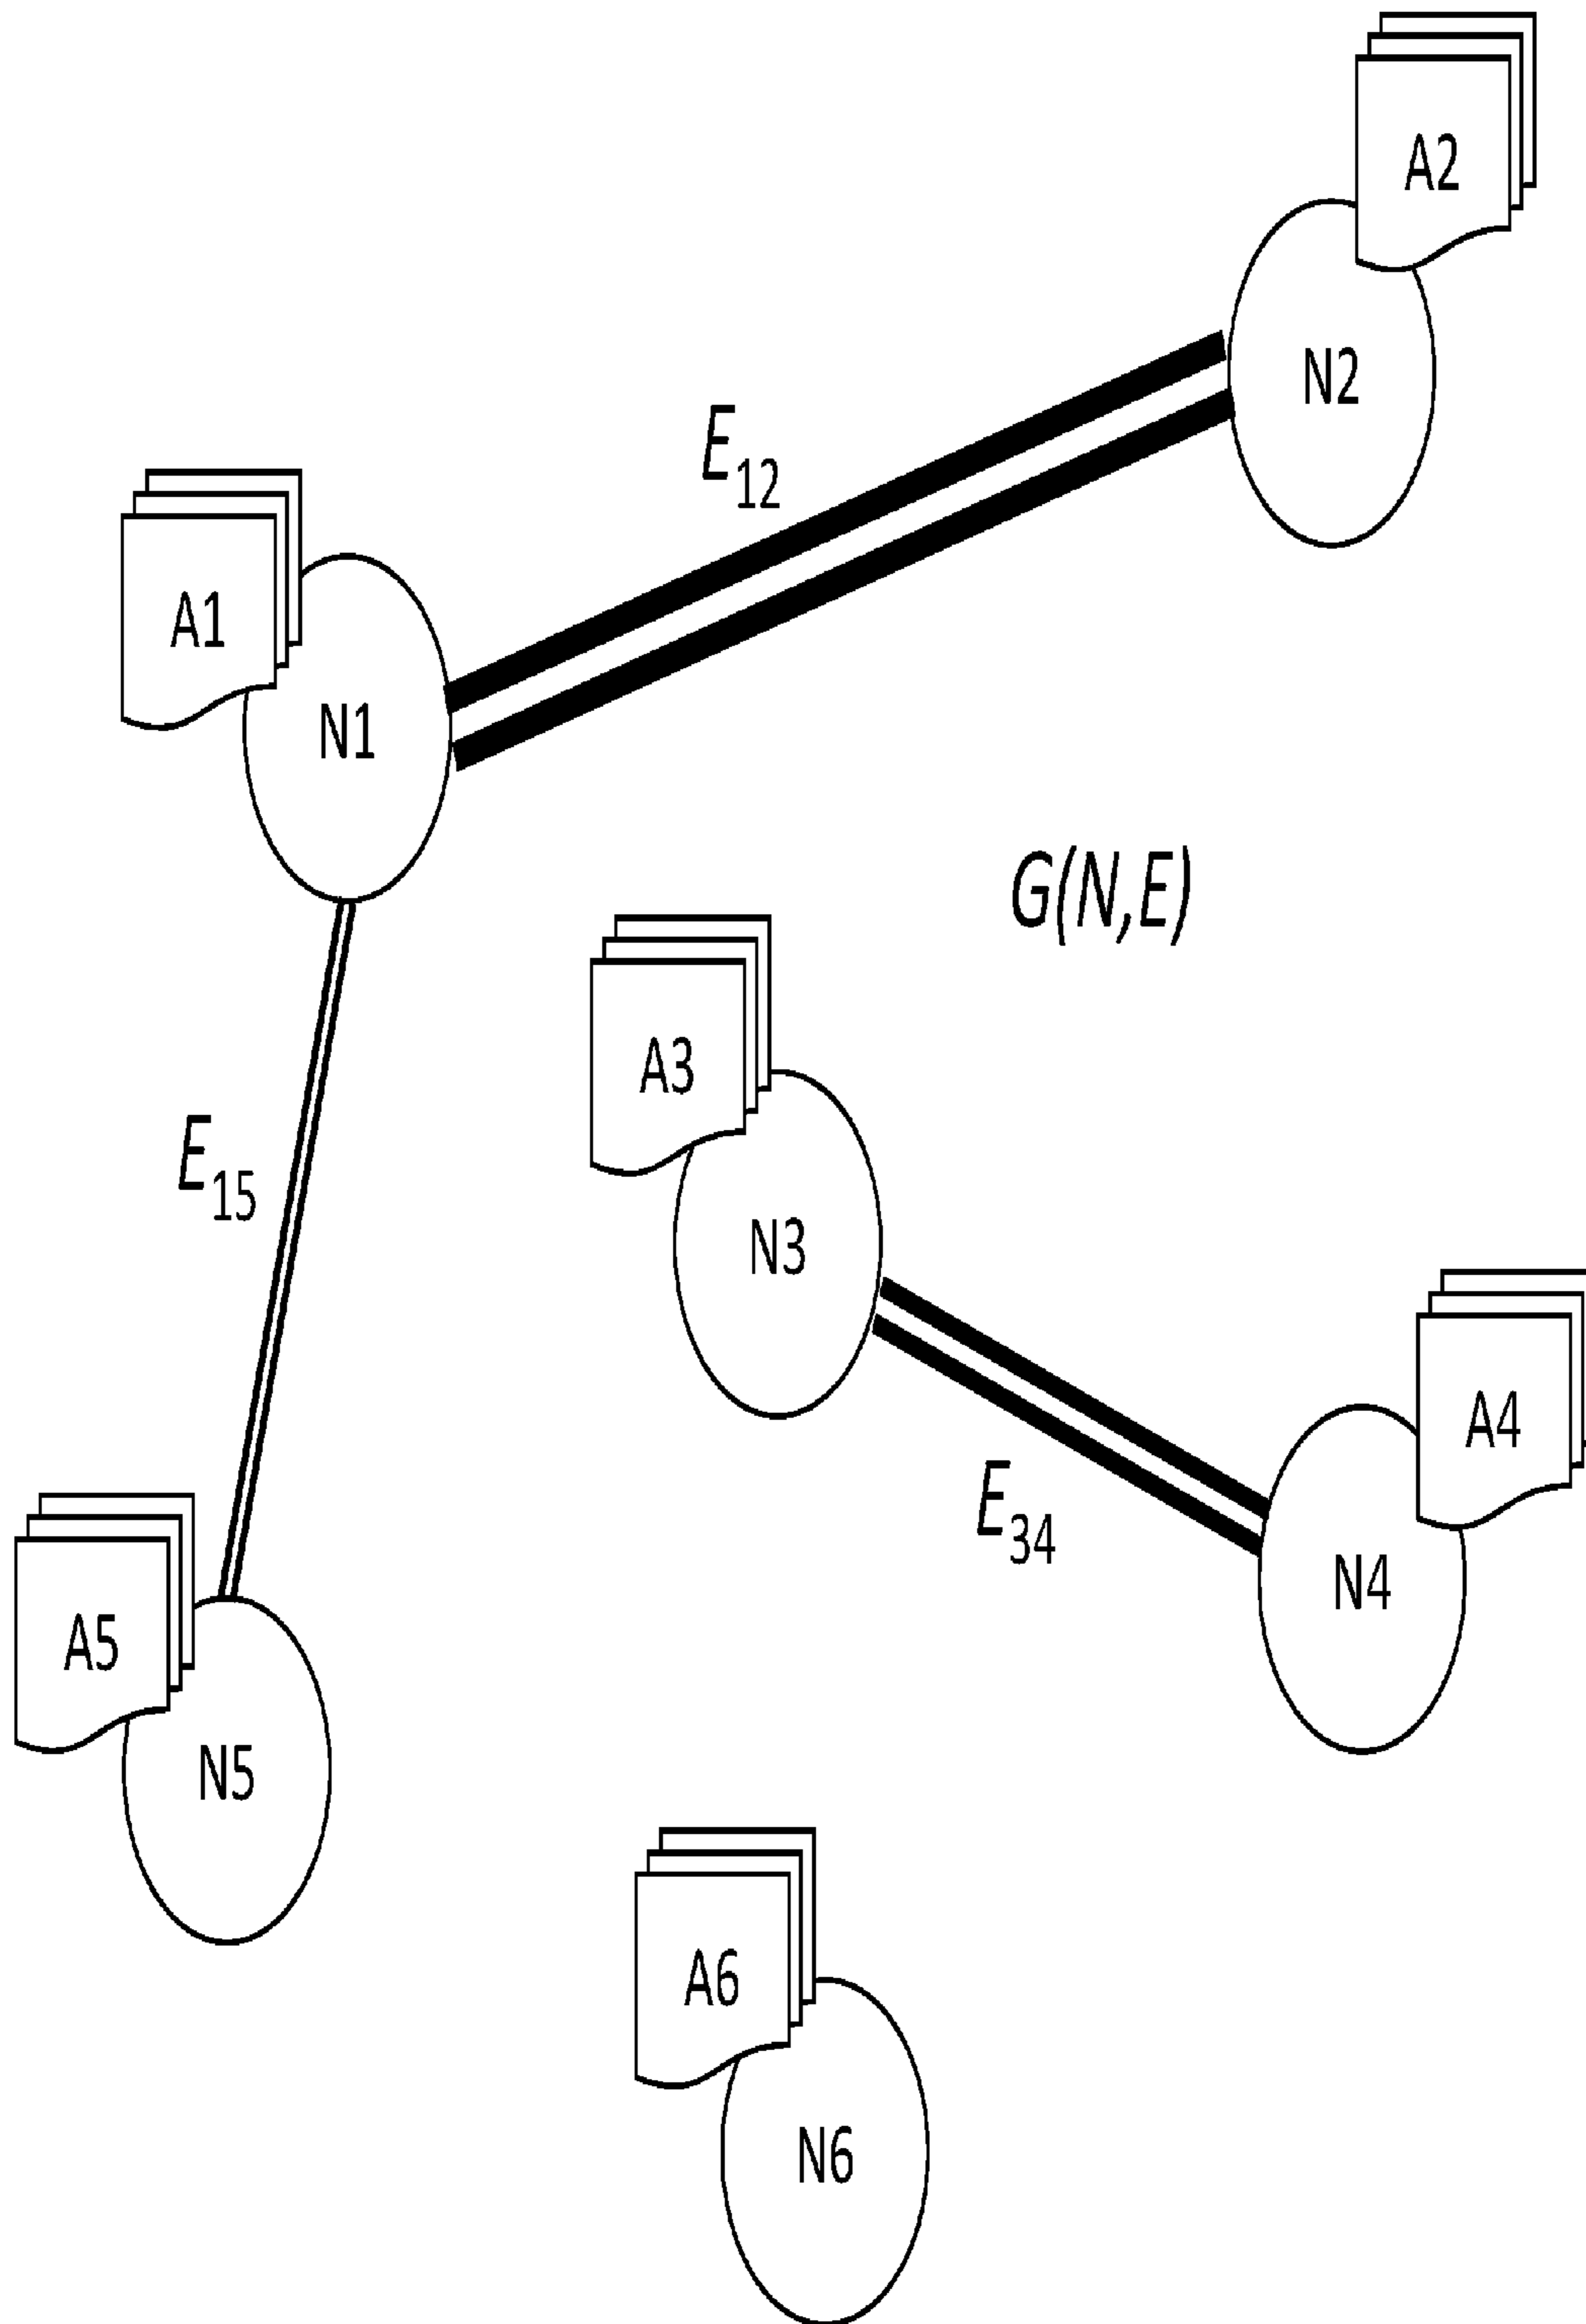


FIG. 1

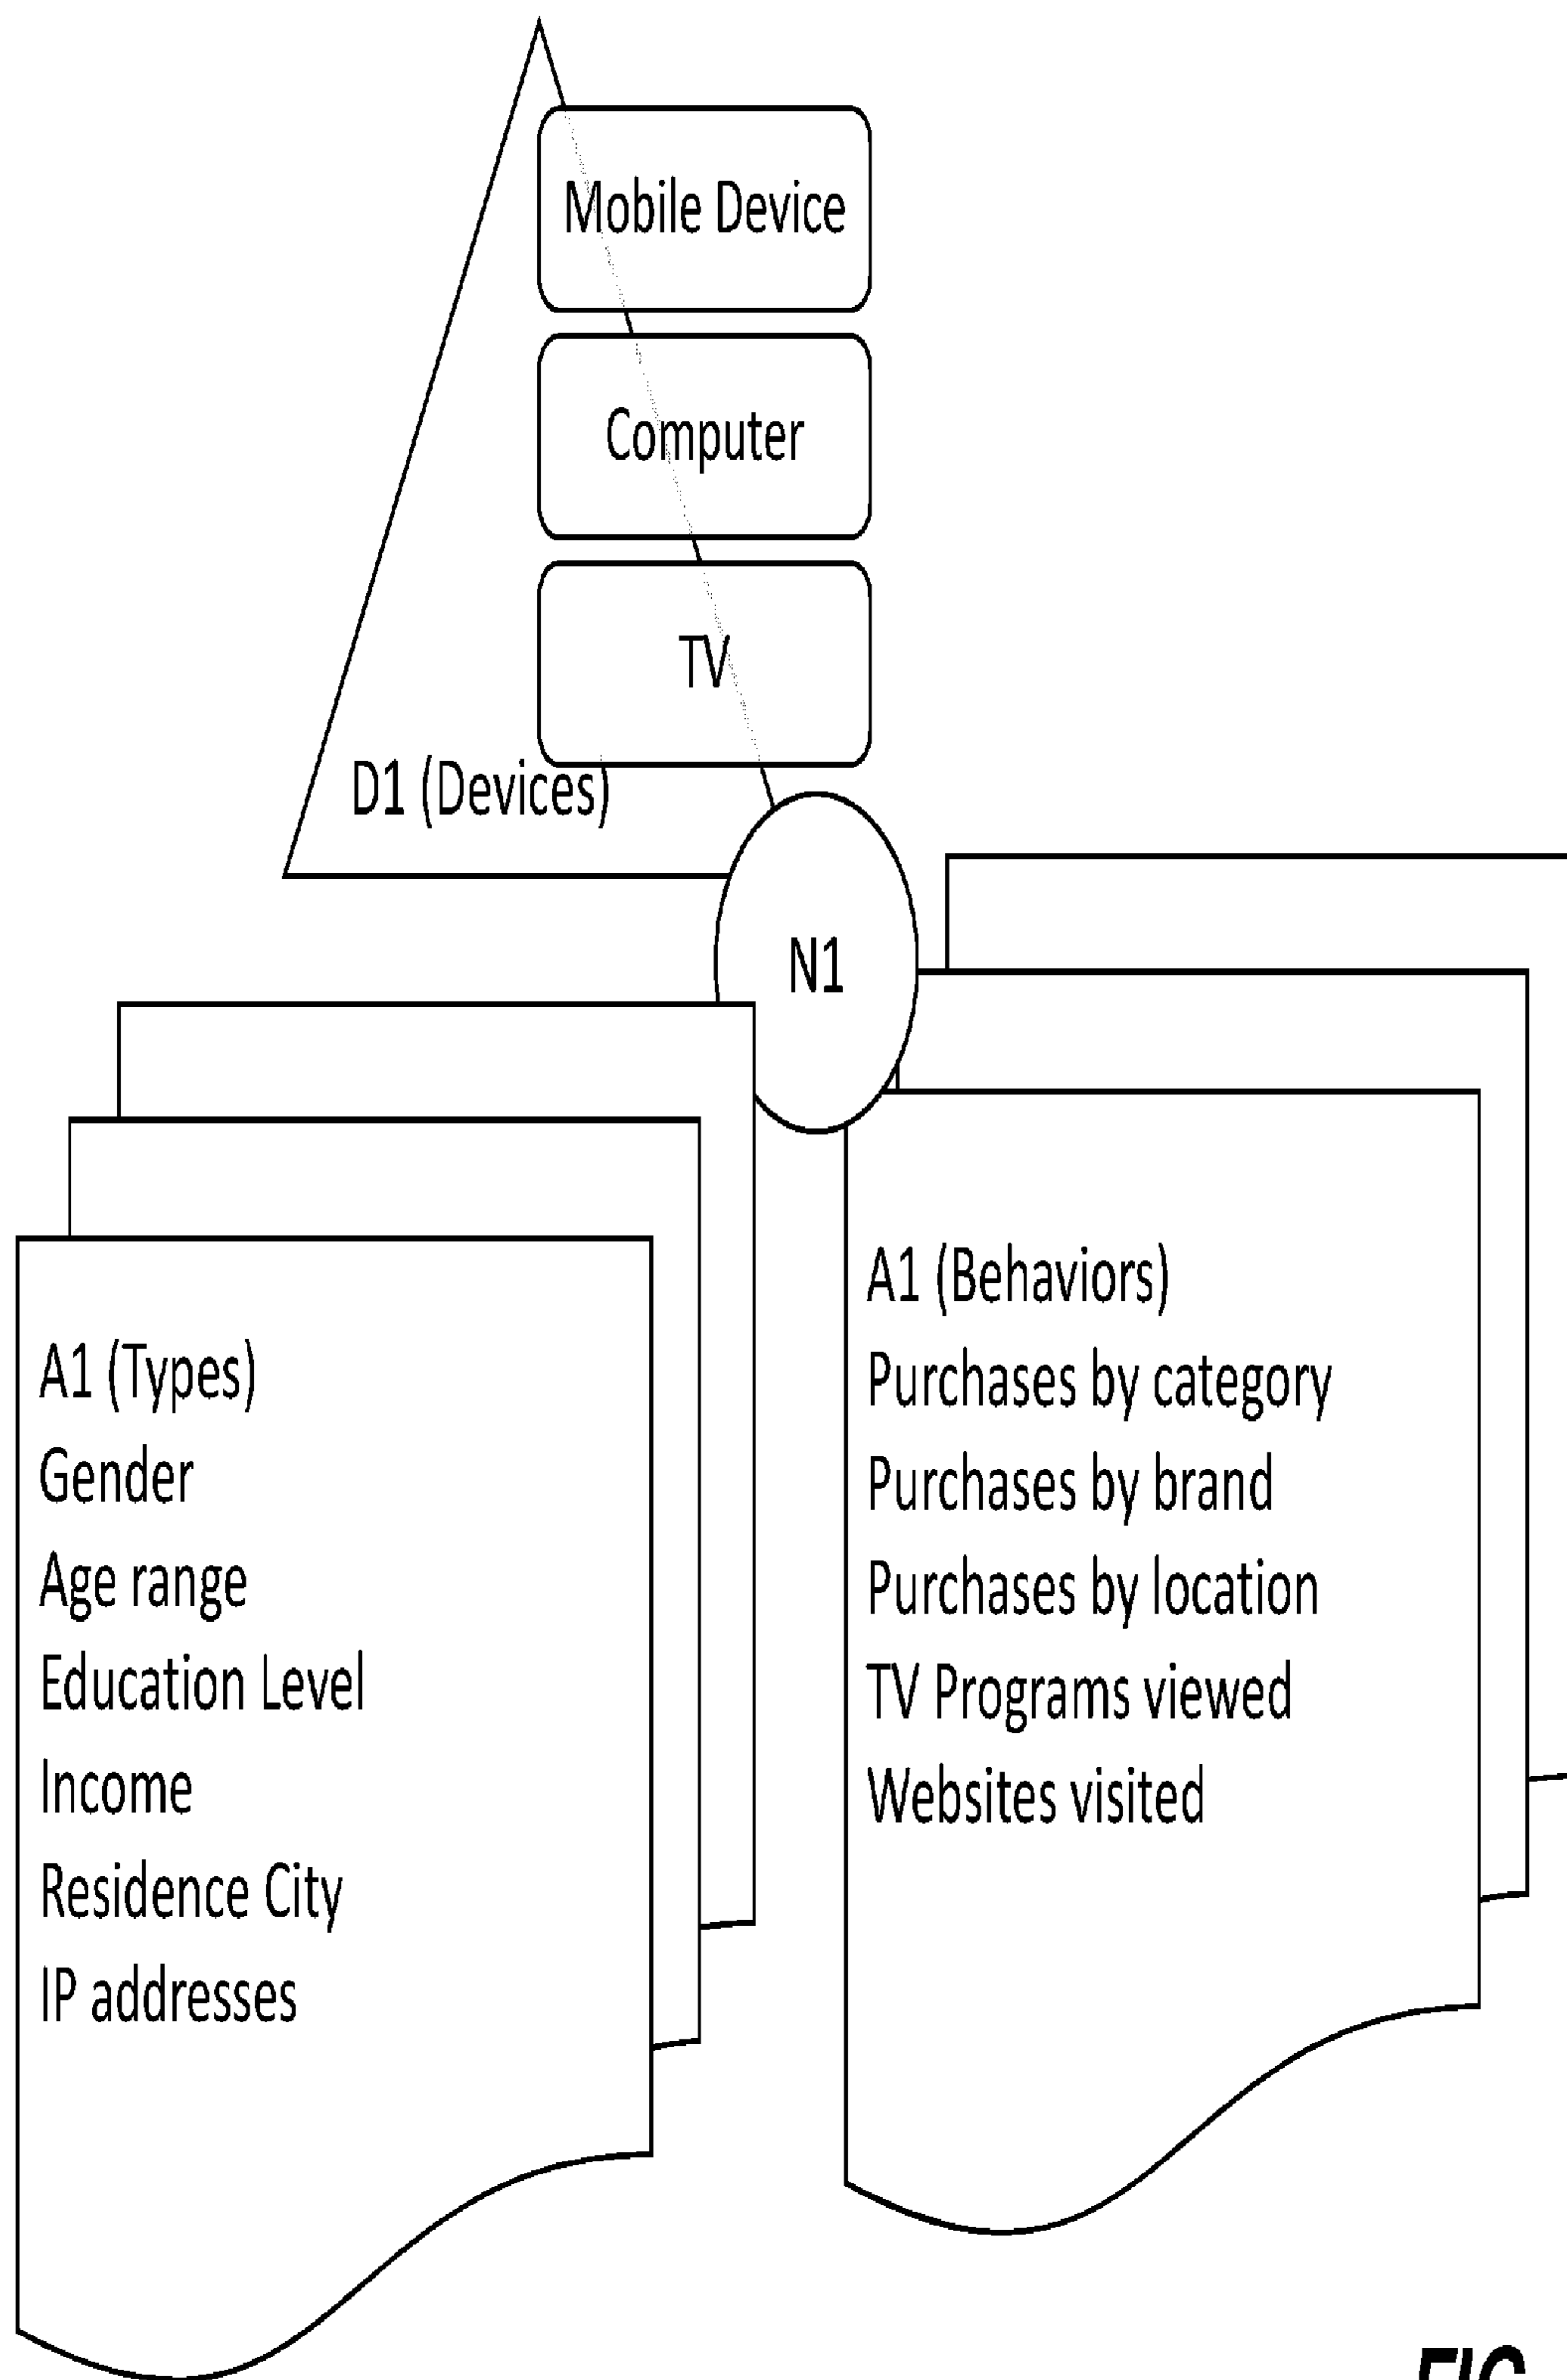


**FIG. 2**

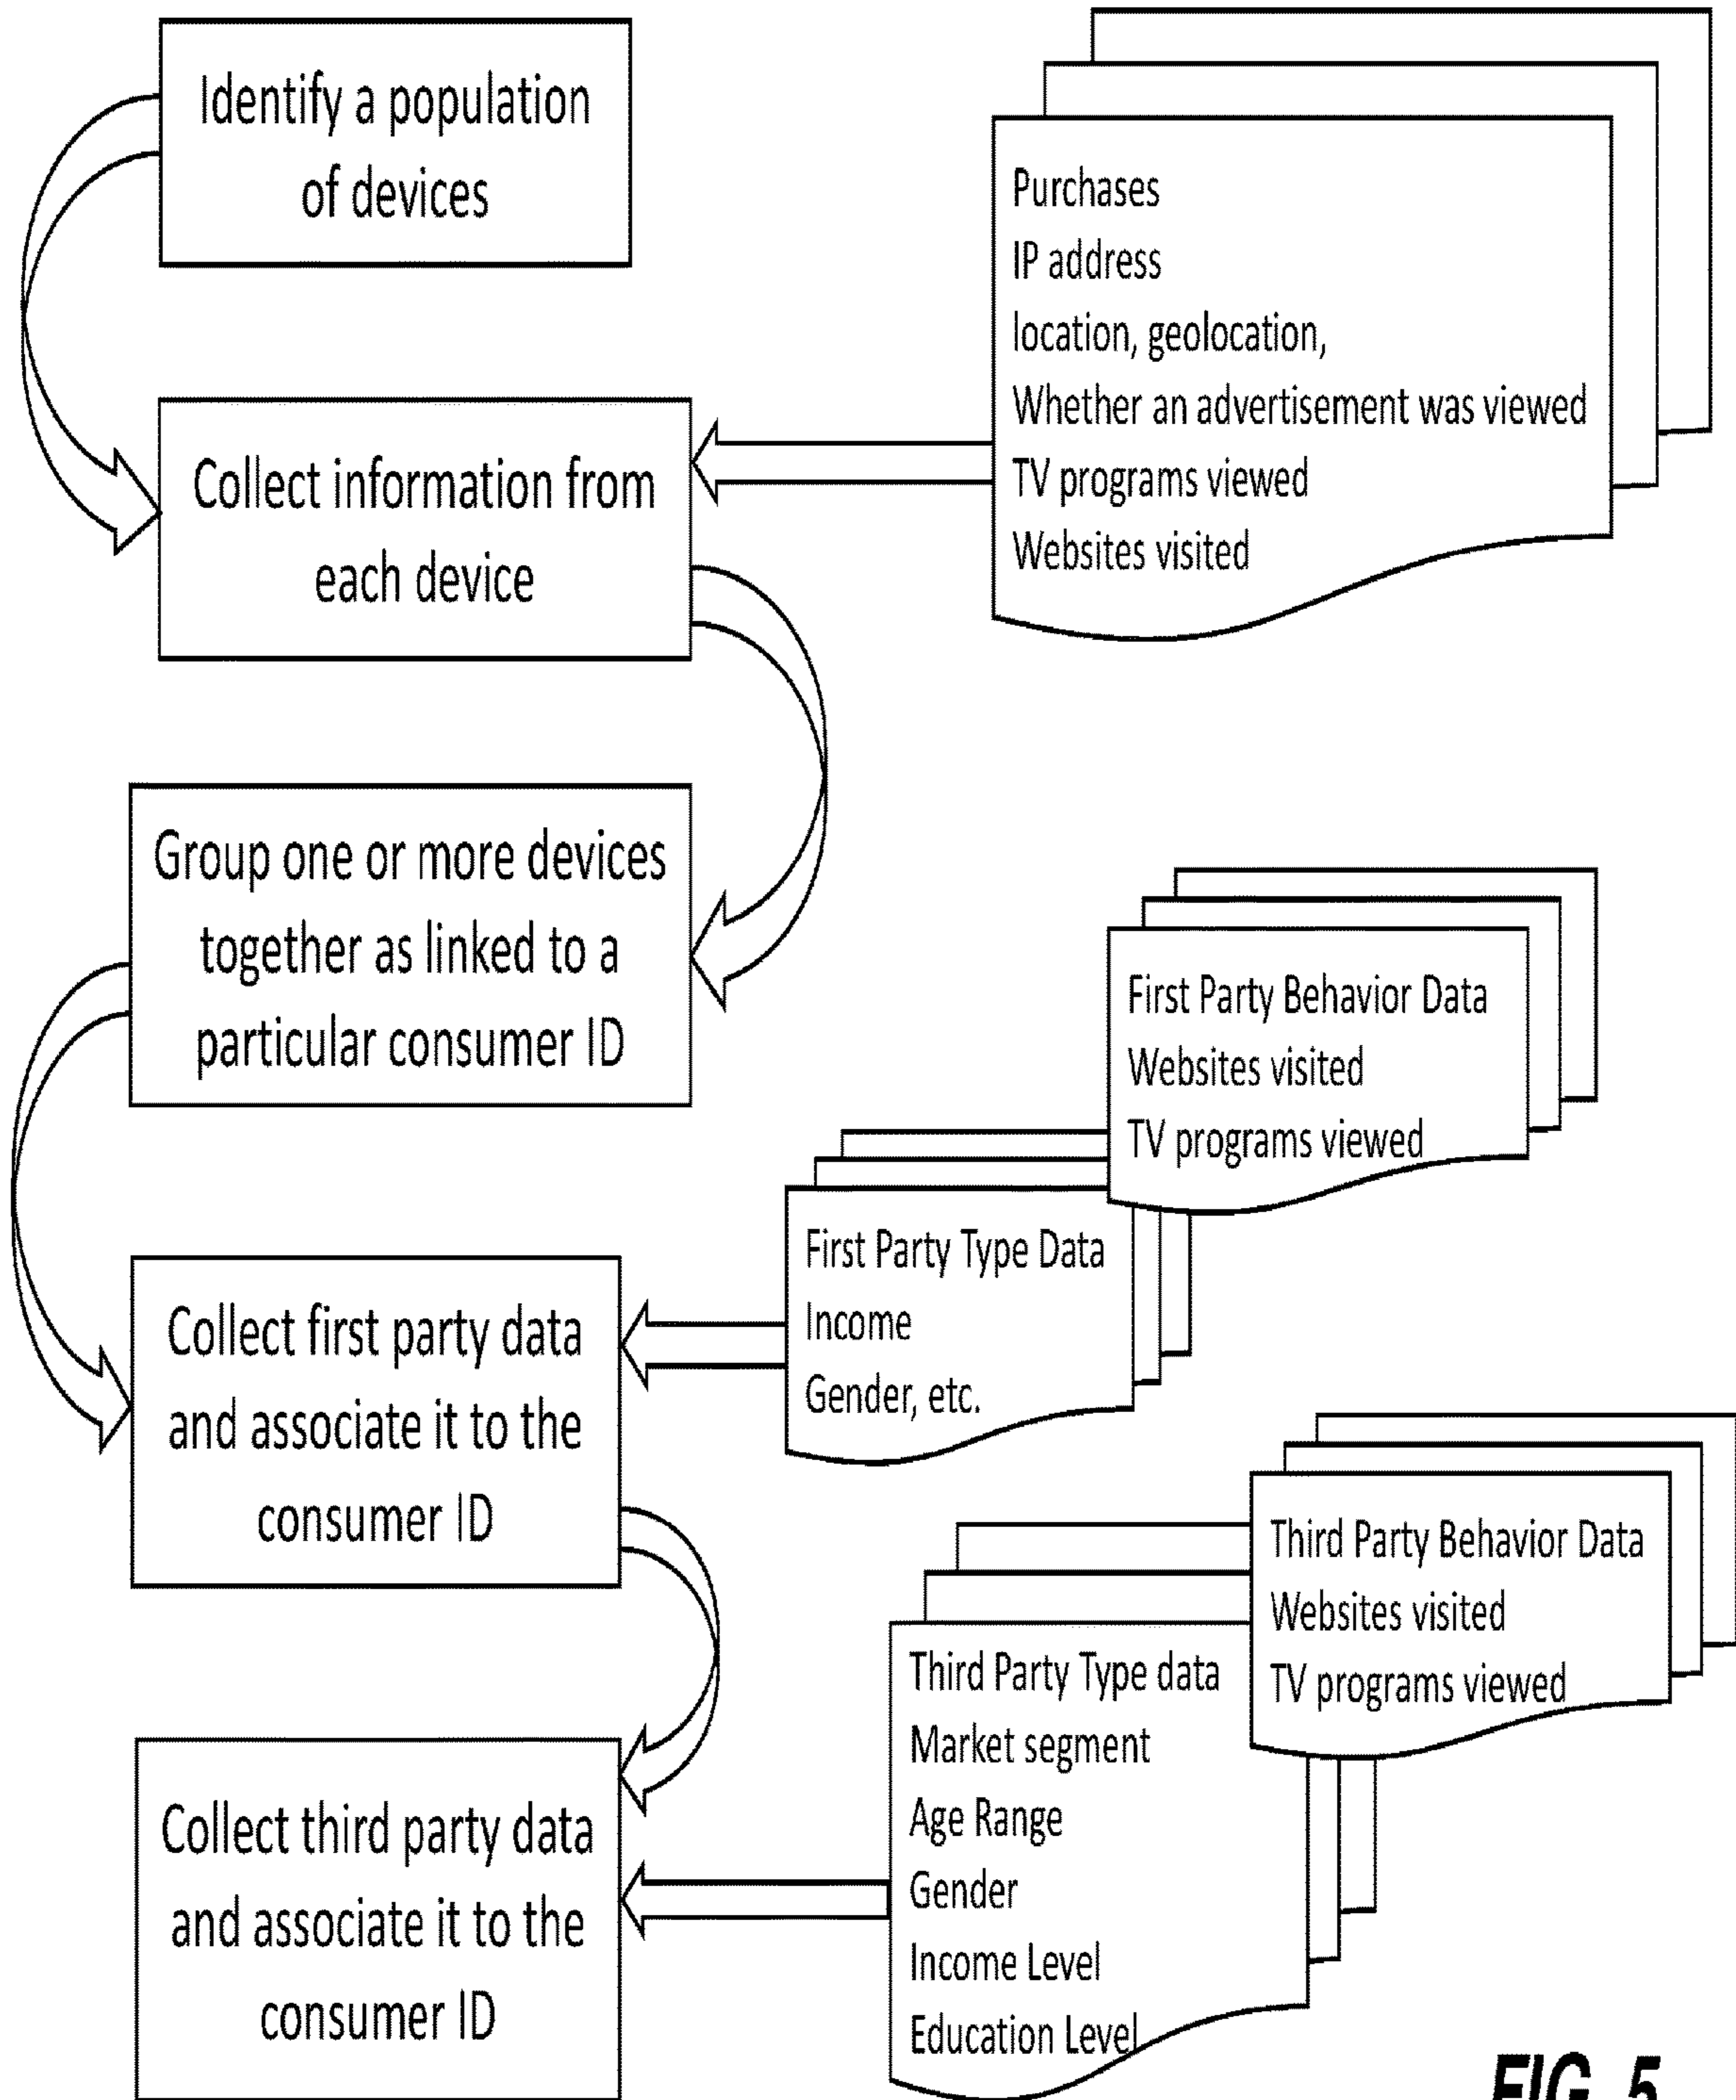


**FIG. 3**

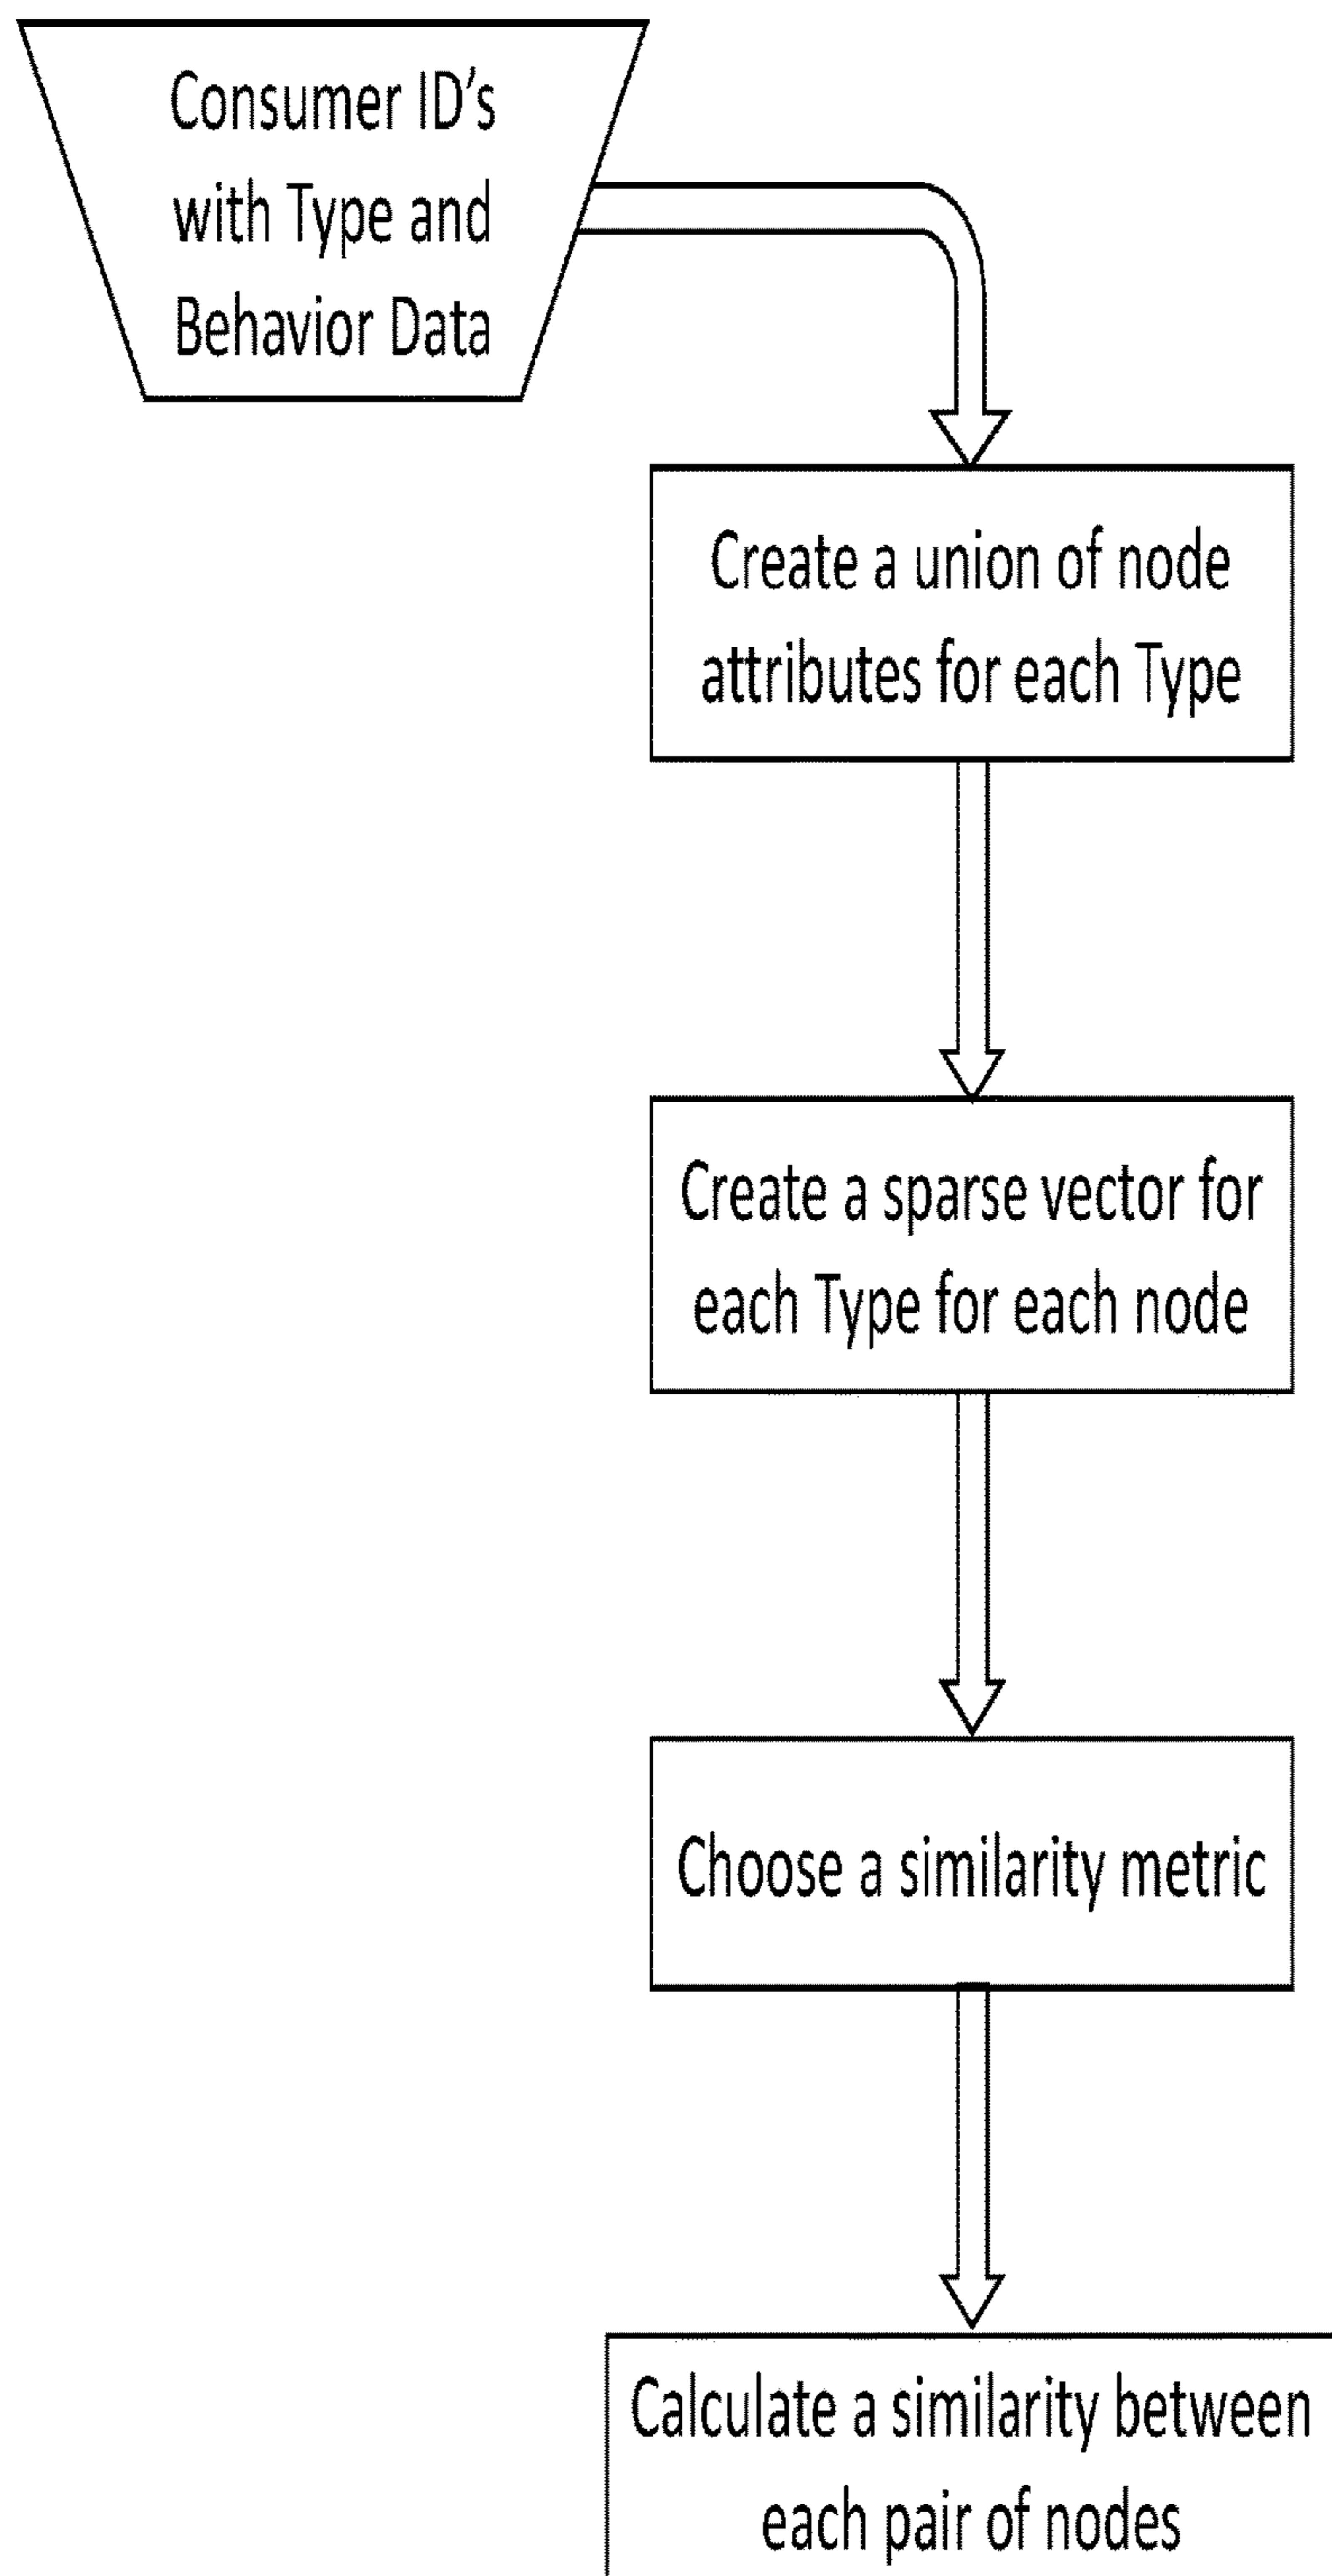




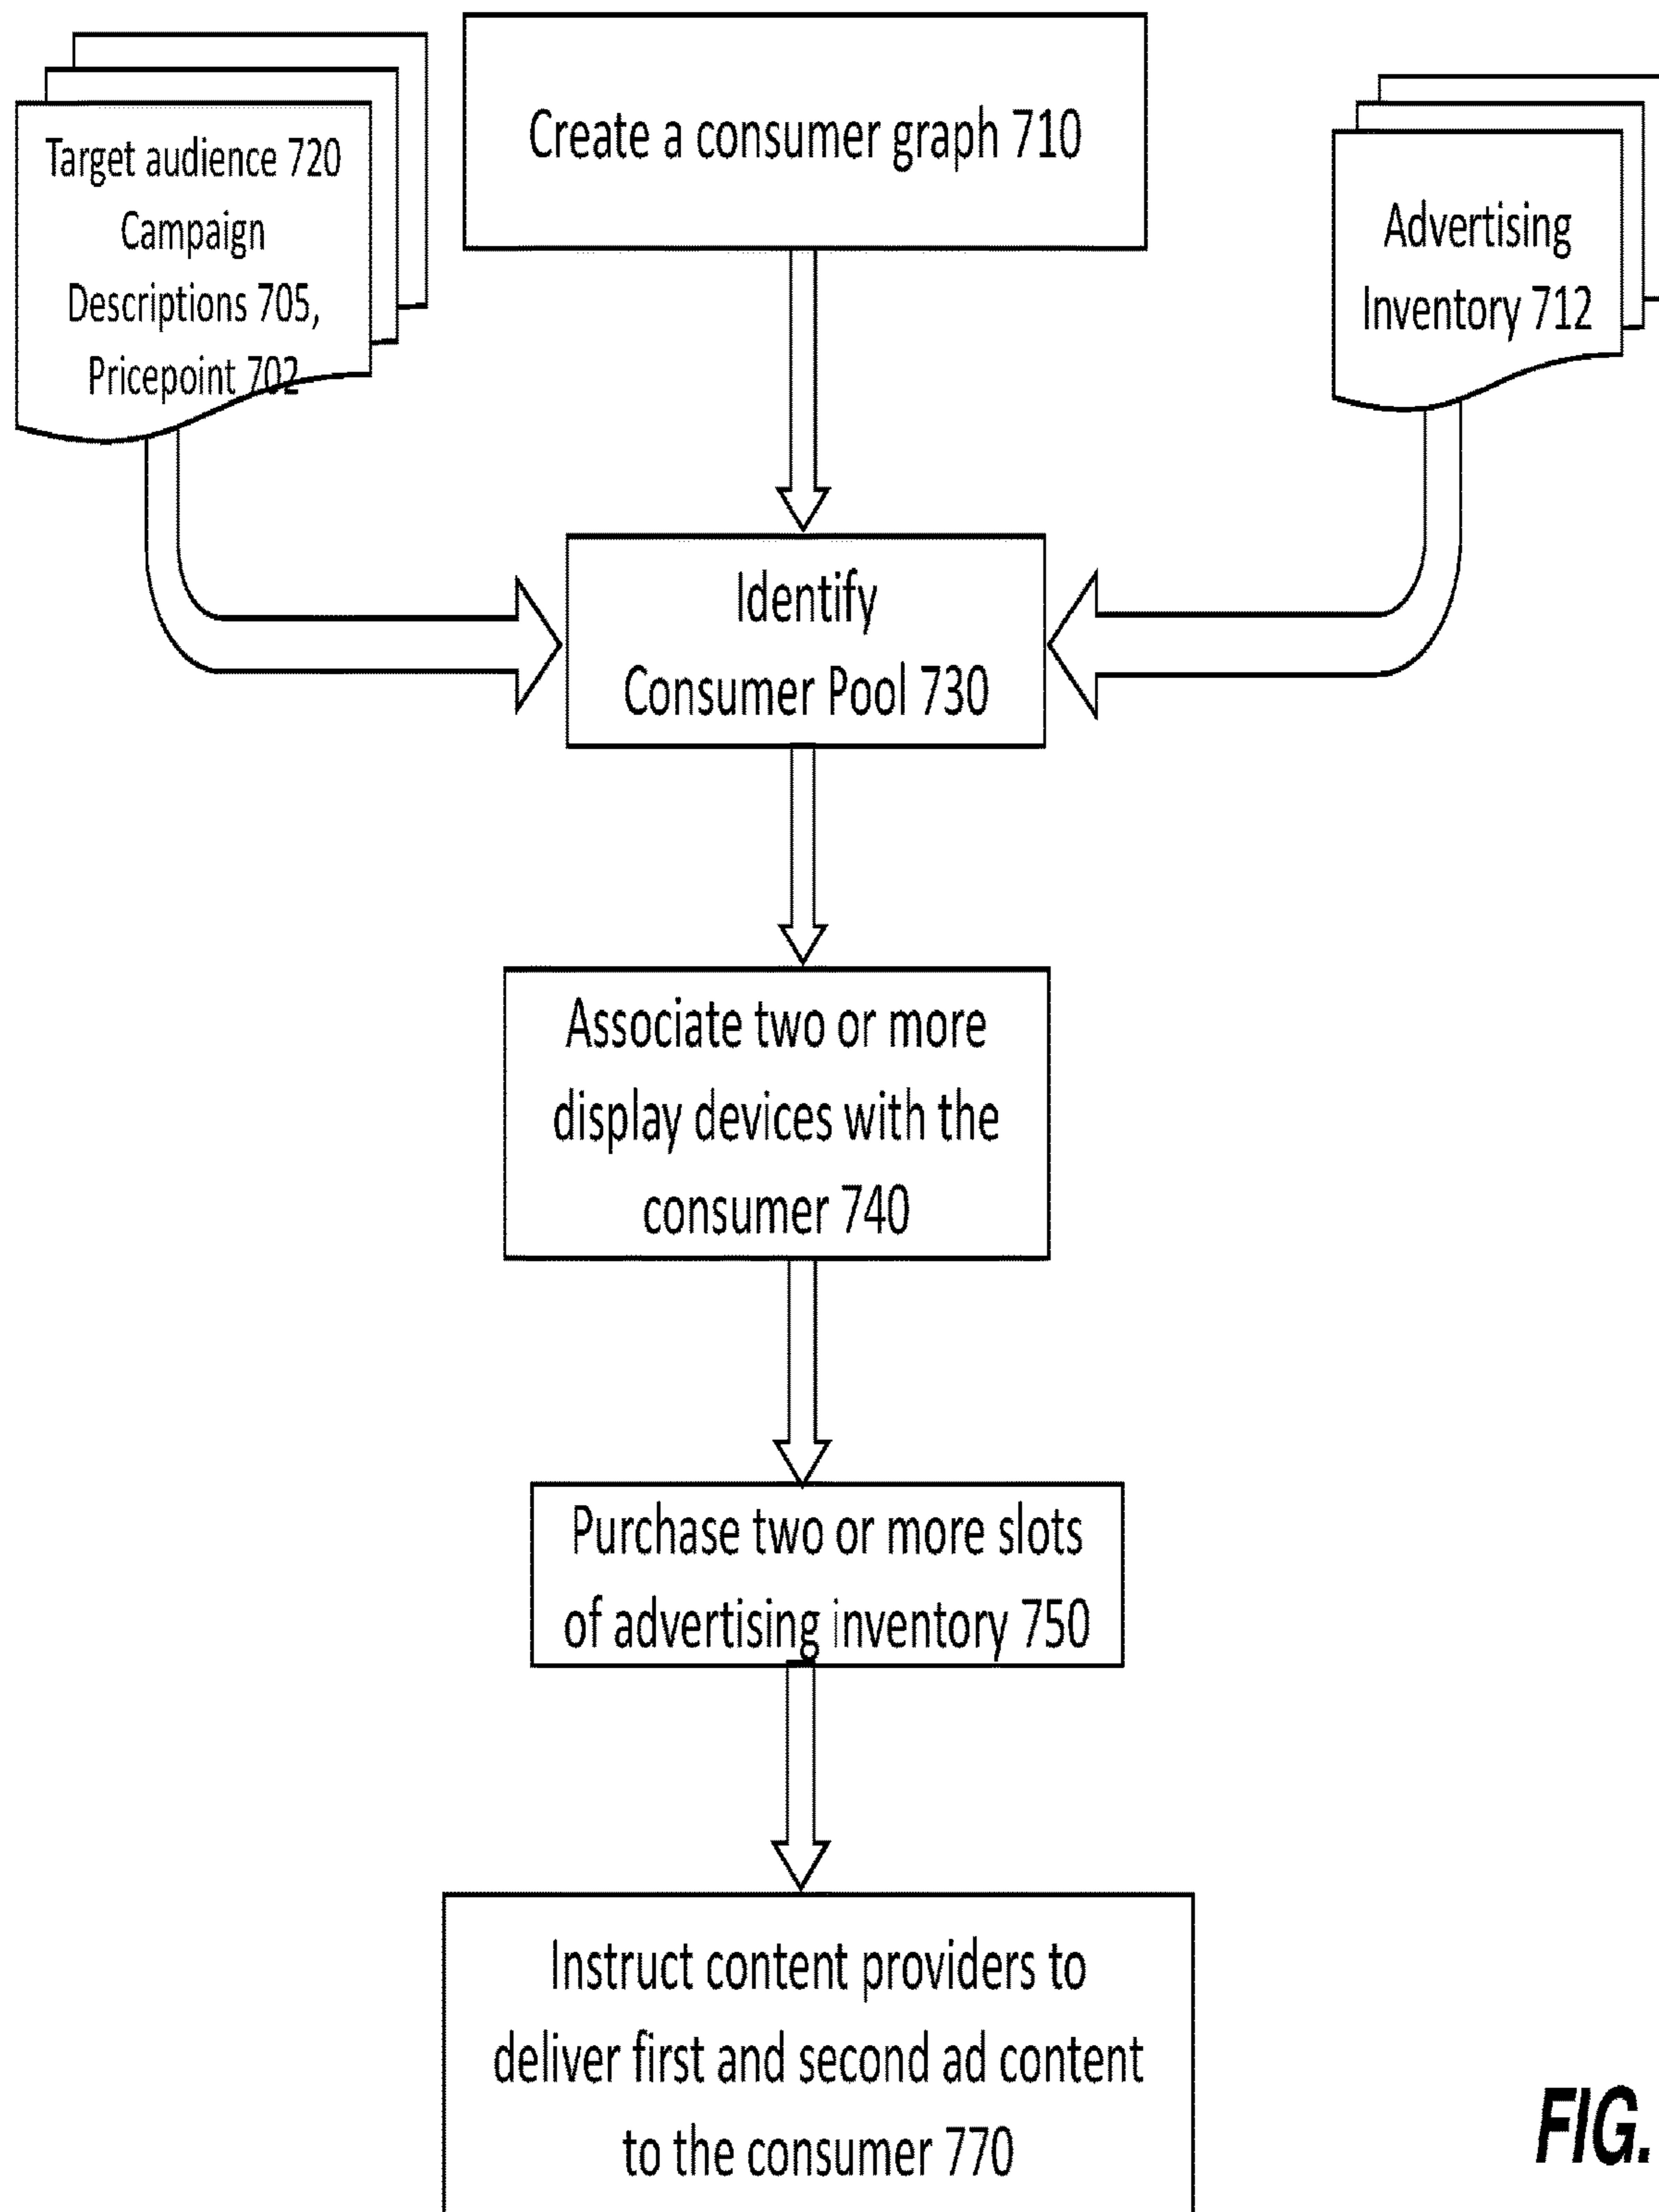
**FIG. 4**



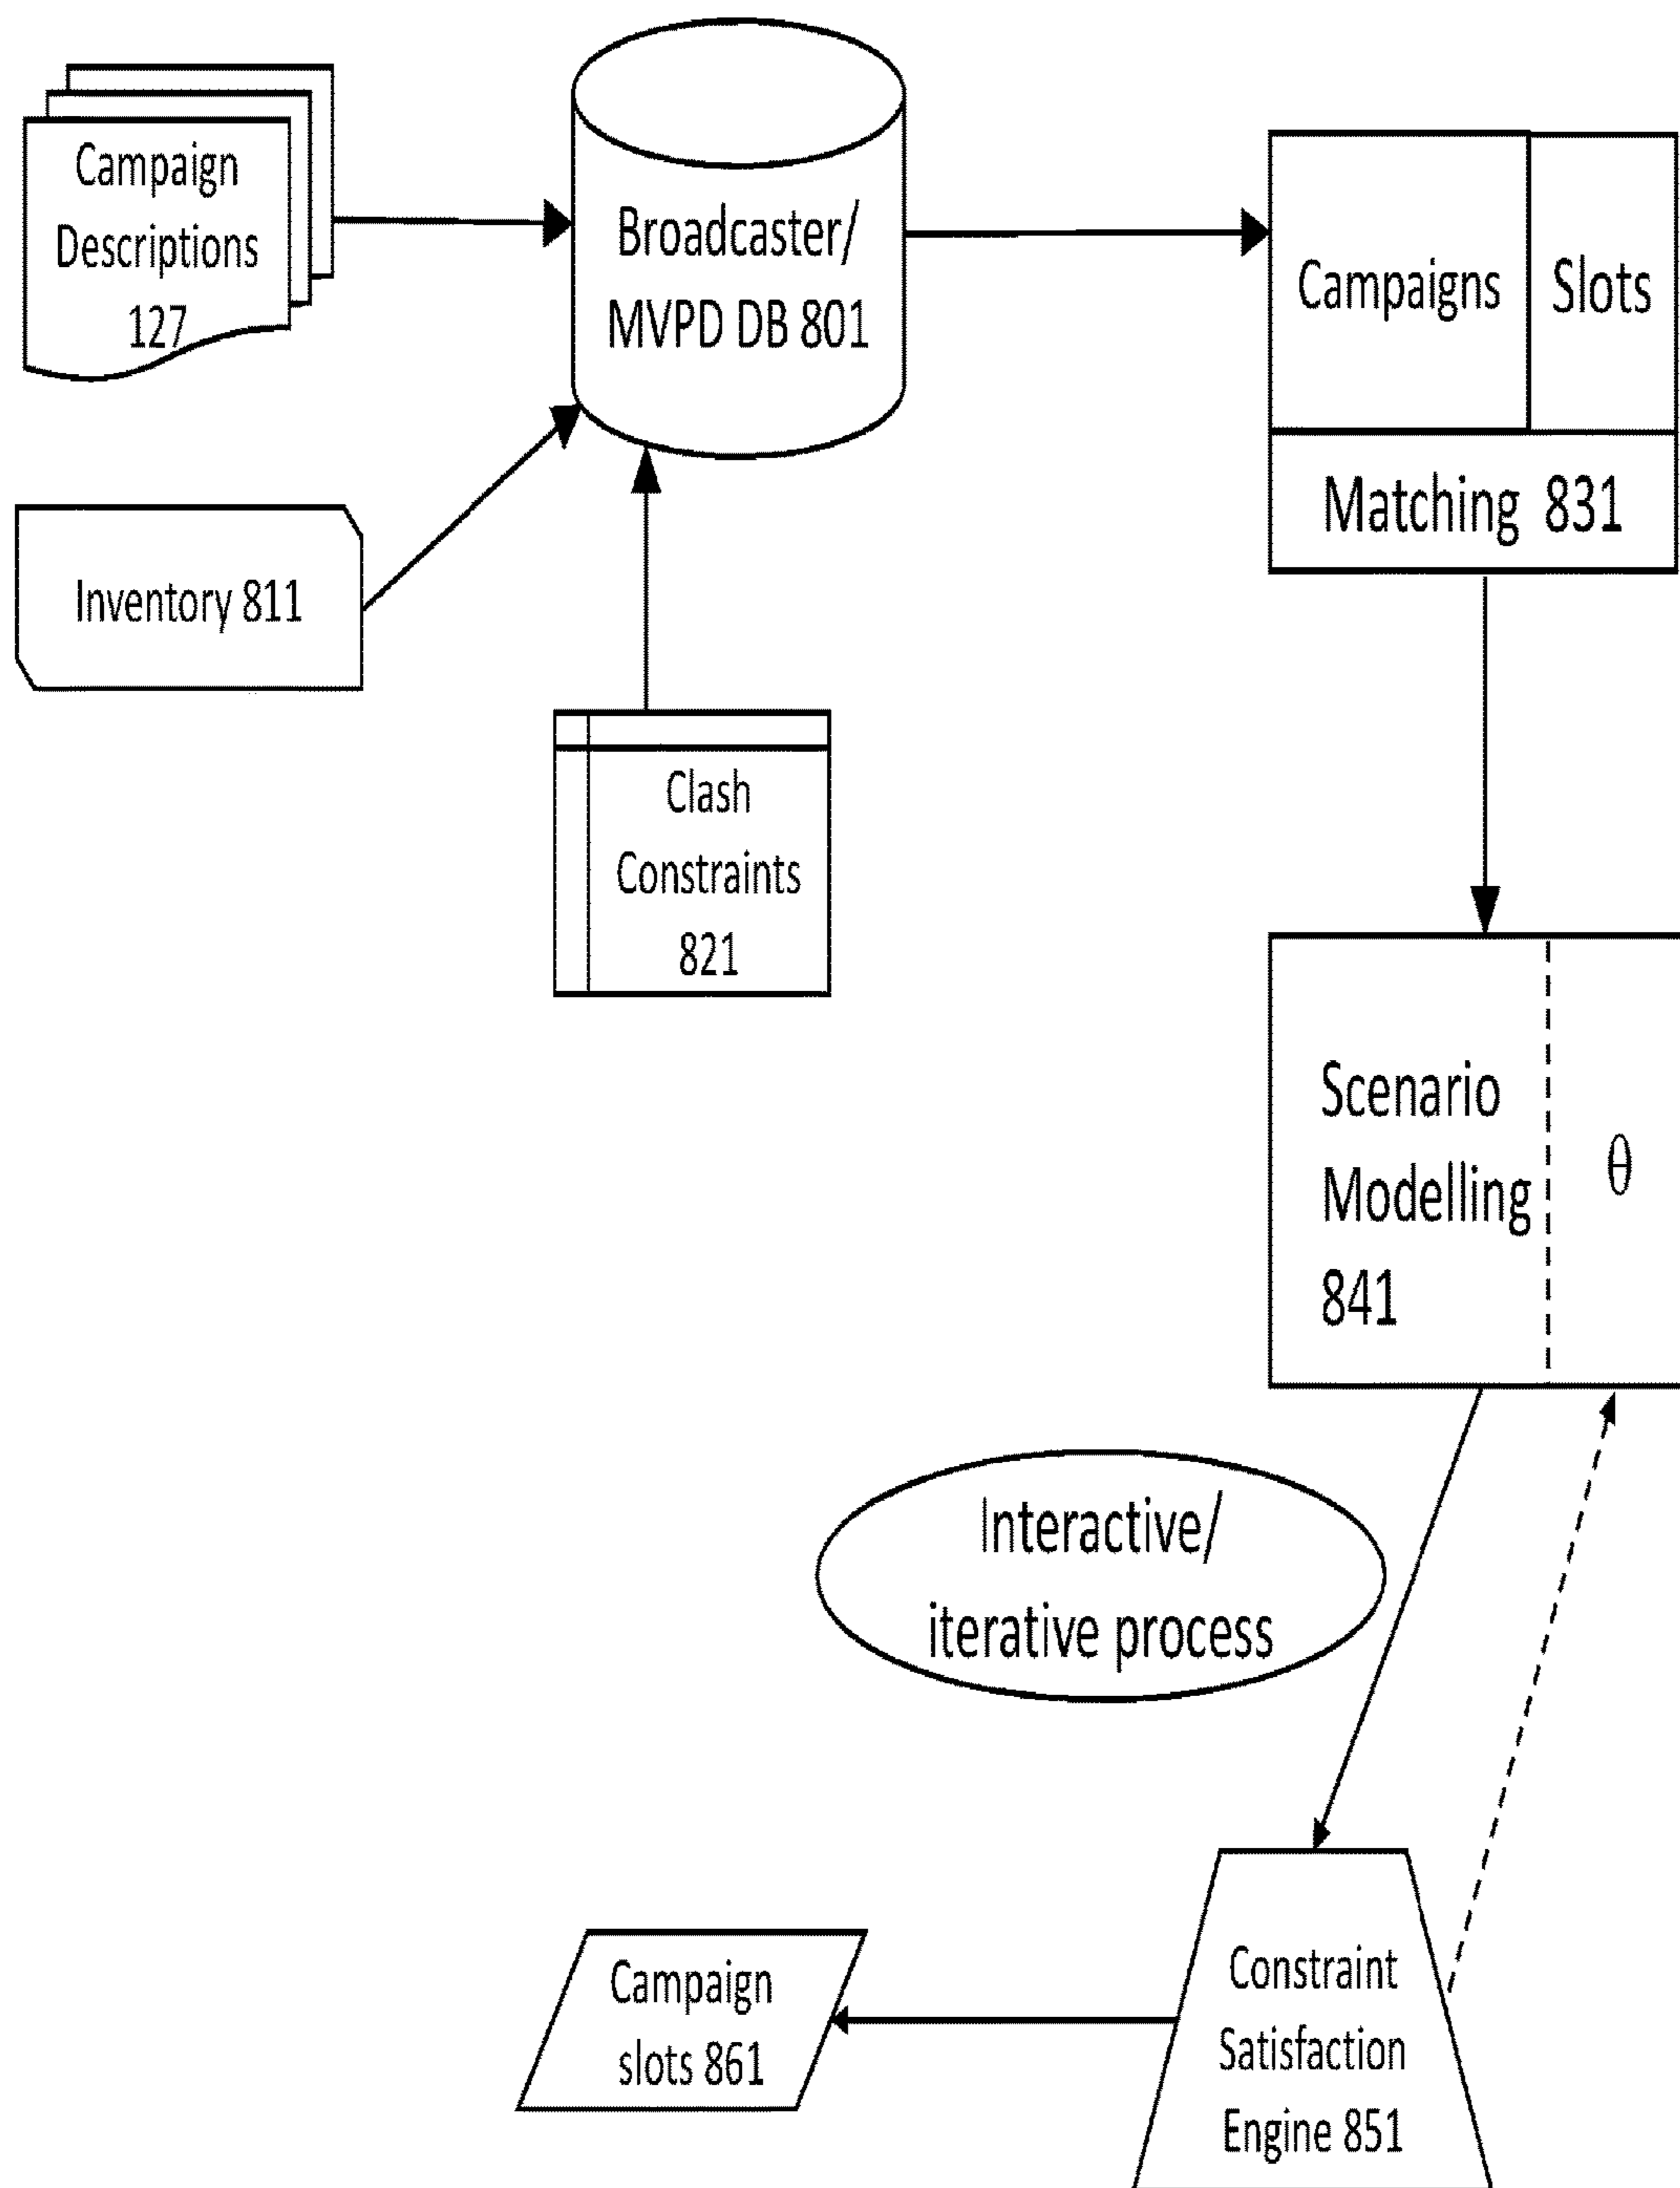
**FIG. 5**



**FIG. 6**



**FIG. 7**



**FIG. 8**


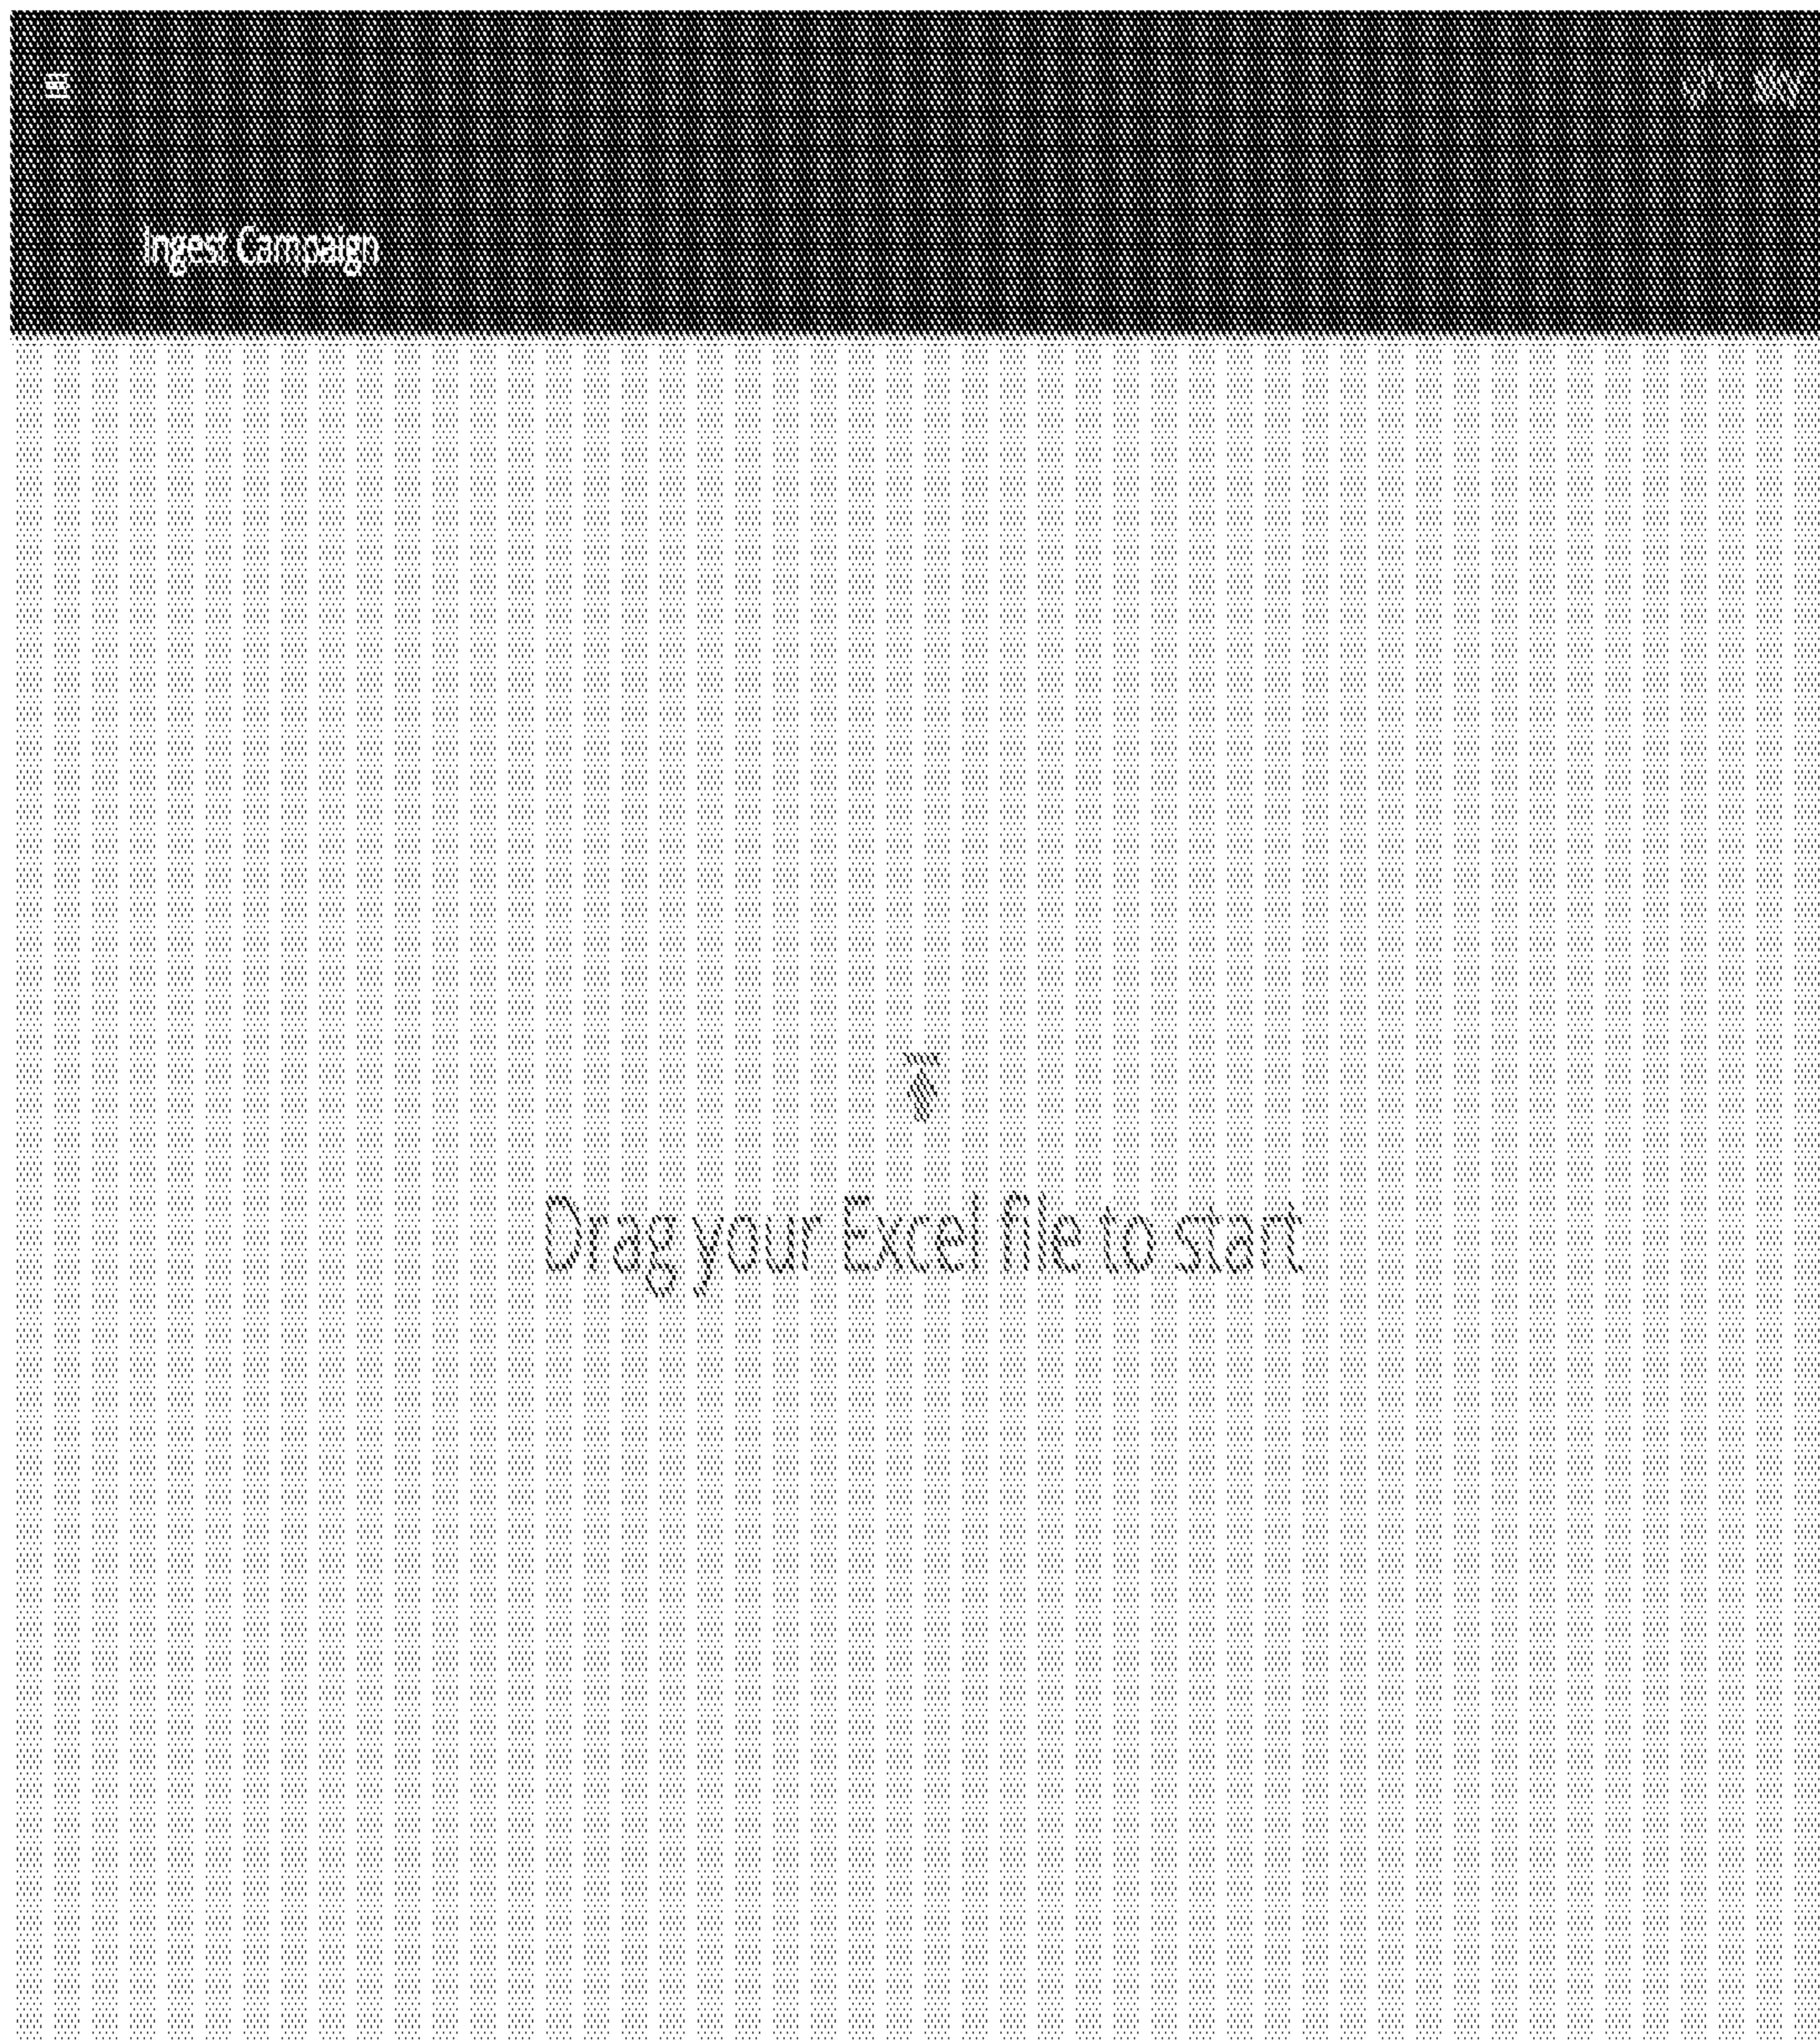


FIG. 9



**FIG. 10A**

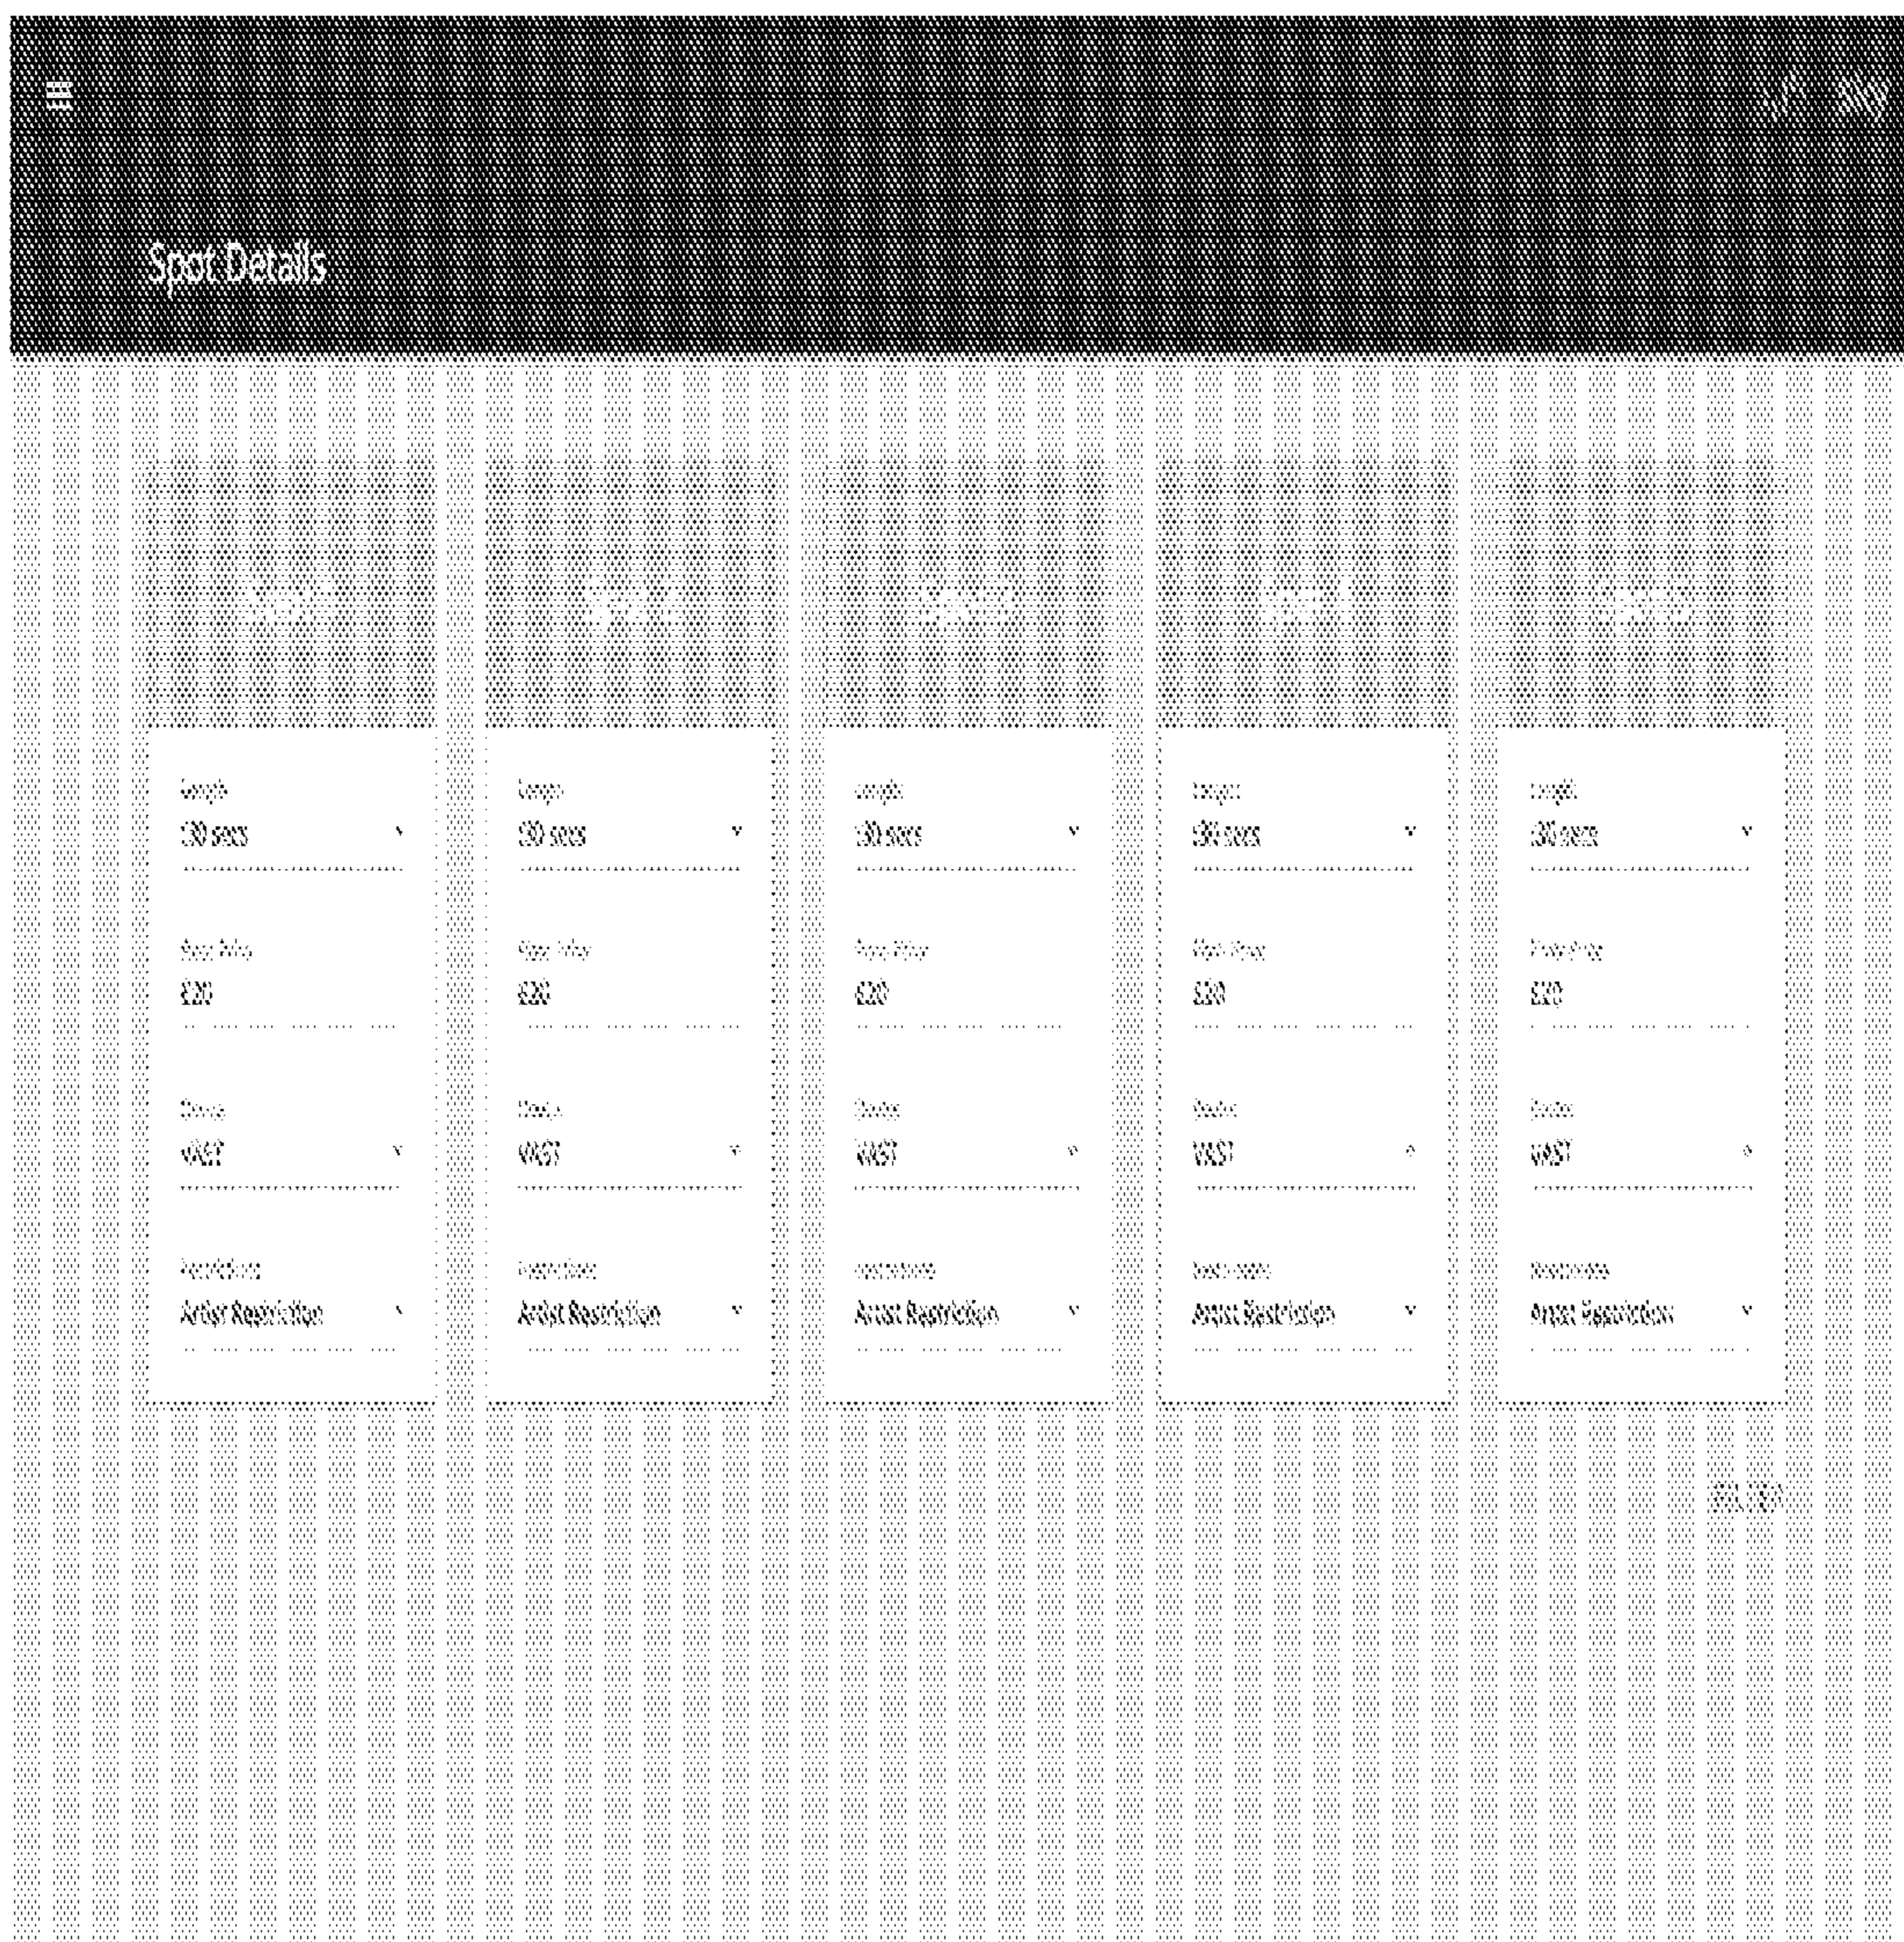


FIG. 10B



Filtered Campaigns

Advertiser	Engagement	Campaign	Date	Description	Exposure	Frequency Cap	CPM	Reach	Campaign ID/Type
Network	30	122	Nov	12345678	4	5	123.54	100%	100%
Advertiser	30	124	Jan	12345678	6	8	123.54	75%	80%
Code	30	125	Feb	12345678	7	9	124.59	80%	70%
Site	30	126	Jan	12345678	5	6	145.33	70%	90%
Client	30	127	Jan	12345678			126.19	70%	70%
Prod	30	128	Feb	12345678			127.74	70%	70%
Advertiser	30	129	Mar	12345678			128.22	80%	80%
Advertiser	30	130	Jan	12345678			129.79	75%	90%

FIG. 10C

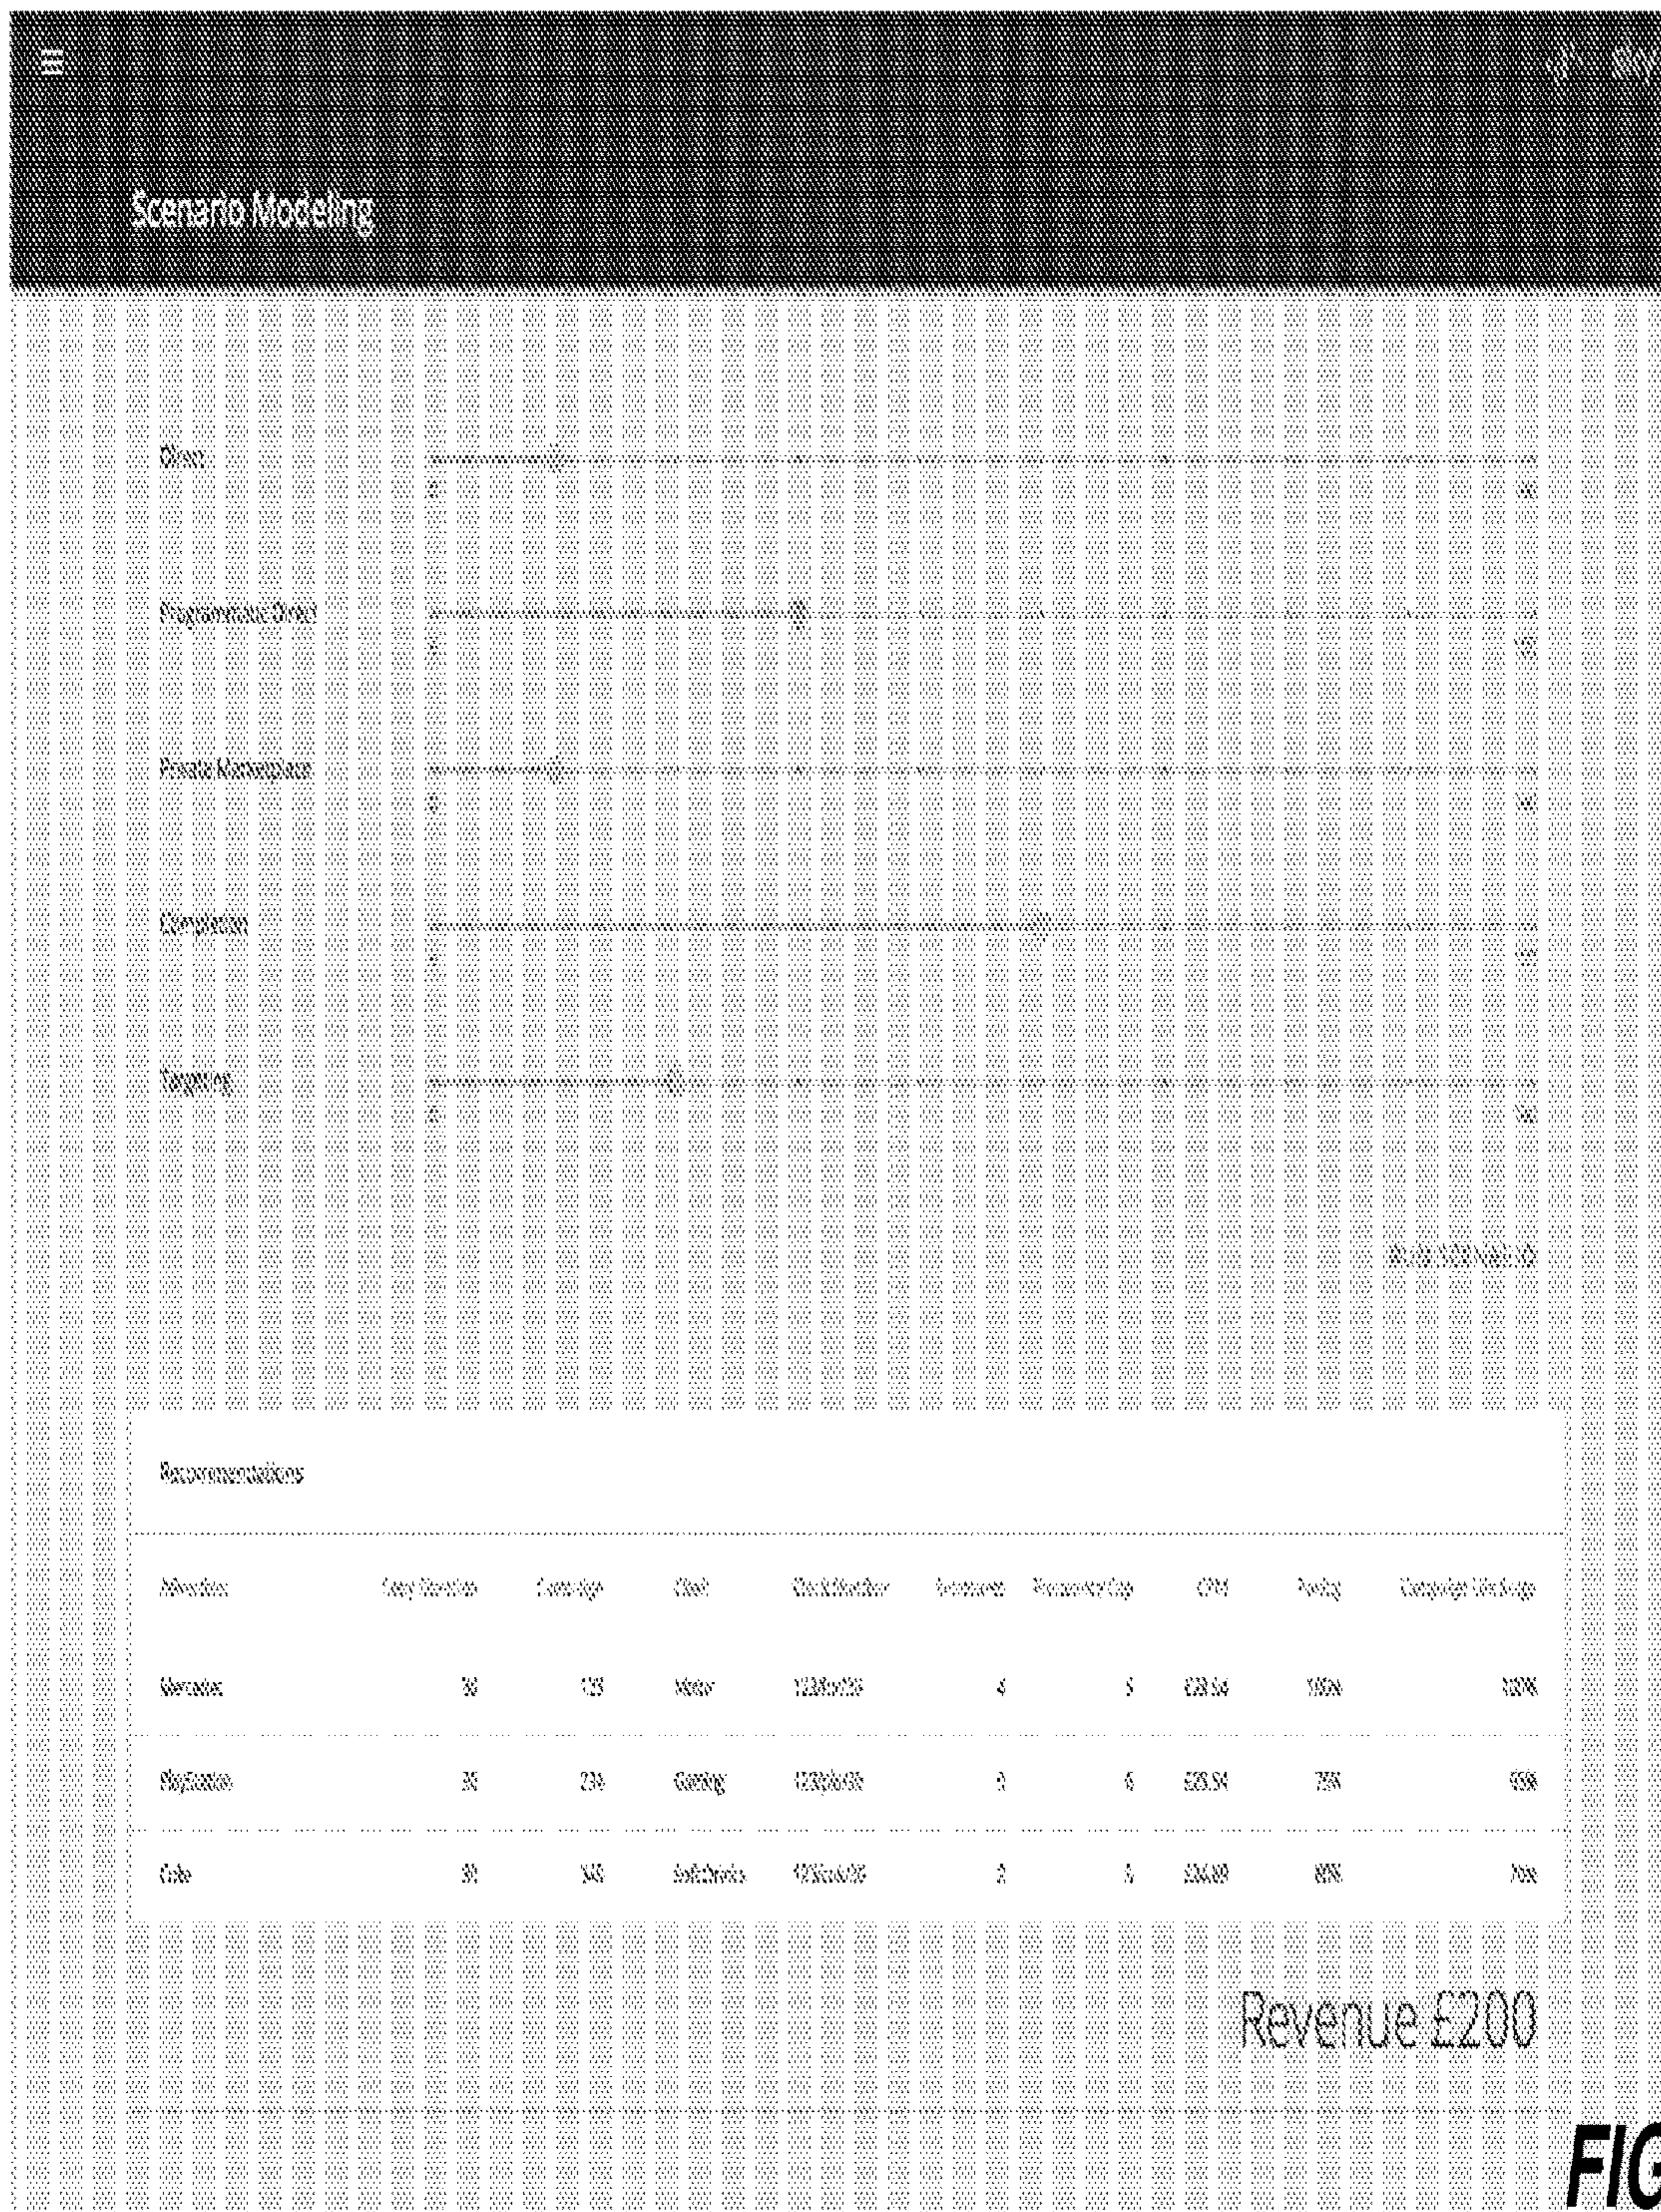


FIG. 10D

## YIELD OPTIMIZATION OF CROSS-SCREEN ADVERTISING PLACEMENT

### CLAIM OF PRIORITY

This application is a continuation of U.S. patent application Ser. No. 15/406,647 filed Jan. 13, 2017, which is a continuation-in-part of U.S. patent application Ser. No. 15/219,262, filed Jul. 25, 2016, which claims the benefit of priority under 35 U.S.C. § 119(e) to U.S. provisional application Ser. No. 62/196,592 filed Jul. 24, 2015 and U.S. provisional application Ser. No. 62/264,764 filed Dec. 8, 2015, U.S. patent application Ser. No. 15/406,647 also claims the benefit of priority under 35 U.S.C. § 119(e) to U.S. provisional application Ser. No. 62/278,888 filed Jan. 14, 2016, and 62/290,387, filed Feb. 2, 2016, all of which are incorporated herein by reference in their entirety.

### RELATED APPLICATIONS

This application is related to U.S. patent application Ser. No. 15/219,259, filed Jul. 25, 2016, entitled “TARGETING TV ADVERTISING SLOTS BASED ON CONSUMER ONLINE BEHAVIOR”, Ser. No. 15/219,268, filed Jul. 25, 2016, entitled “CROSS-SCREEN MEASUREMENT ACCURACY IN ADVERTISING PERFORMANCE”, Ser. No. 15/219,264, filed Jul. 25, 2016, entitled “SEQUENTIAL DELIVERY OF ADVERTISING CONTENT ACROSS MEDIA DEVICES”, and to provisional application Ser. No. 62/317,440, filed Apr. 2, 2016, all of which are incorporated herein by reference in their entirety.

### TECHNICAL FIELD

The technology described herein generally relates to improving and managing a cross-screen advertising strategy for advertisers, and more particularly relates to a system and method for targeting advertising content to consumers on TV and mobile devices that can be operated from within a multi-channel video programming distributor’s environment.

### BACKGROUND

Video advertisements are among the most advanced, complicated, and expensive, forms of advertising content, particularly those that are intended to be delivered to the growing number of virtual reality devices as well as those that demand some level of user-interactivity to achieve their intended effect. Beyond the costs to produce video content itself, the expense of delivering video content over the broadcast and cable networks remains considerable, in part because television (TV) slots are premium advertising space in today’s economy. Furthermore, TV is no longer a monolithic segment of the media market. Consumers can now spread their viewing of video content, particularly premium content, across TV, DVR, and a menagerie of over-the-top and on-demand video services viewed across smart TVs, gaming consoles, and mobile devices, as well as traditional TVs.

In short, TV viewing is transforming to digitally distributed viewing, as audiences watch proportionately less live broadcasting and more in a video on demand (VOD) or streaming video format. Correspondingly, broadcasters are becoming able to identify specific market segments with increasingly greater precision, and therefore can offer adver-

tisers who purchase time on their channels the opportunity to target their content to highly relevant slices of the population.

Adding online consumption to the list of options available to any given consumer, only leads to greater complexity in the process of coordinating delivery of video adverts to a relevant segment of the public, due at least in part to the number of different types of devices that can be online today—and in the future—as well as the fact that a given individual may interact with many different types of online devices in any one day. This complexity means that the task of optimizing delivery of advertising content today far exceeds what has traditionally been necessary, and what has previously been within the capability of experienced persons. The data needed to fully understand a given consumer is fragmented as each individual and household views more and more media in a disparate fashion by accessing a network of devices.

Furthermore, advertising scheduling problems grow combinatorially as the number of advertisements and inventory slots increase. As such, existing methods for preventing the accidental breach of a constraint (i.e. media content regulation or contractual duty) on advertising placements are highly ineffective for even modest volumes of constraints. While existing advertising yield optimization software tools that allow for simple relational database exploration of constraints, they are not able to process the analytics and produce suggestions for advertising placement based on those constraints.

For many companies, the analytical work that goes into developing an advertising strategy today still requires manual contributions by human analysts. This is especially the case for low volume purchasers of advertising inventory. Advertising strategies are also generally fixed, meaning approaches to advertising strategy are conditioned on certain assumptions that are inflexible and limited to what manual processes are able to achieve.

It is also challenging to assign numeric values to soft constraints such as business goals and internal seller-side priorities for existing advertiser-customers. Systems and methods are not able to explore a wide range of weighted approaches to selling advertising inventory because doing so is overly time consuming and expensive because it requires human modeling on Excel-style spreadsheets. Also, in the current market where new demand sources can be added in real time, this complexity of new potential campaigns can’t be accommodated by manual or spreadsheet approaches. The market is real time and always on.

The current state of advertising strategy is analogous to where financial trading was before the creation of financial strategy tools such as E-Trade, which facilitates automated buying, and financial advisors such as Fidelity for investment planning.

In today’s advertising strategy, human analysts guide the selection of advertising inventory based on, for example, Excel data tables and other static data management tools. This results in inefficient selection of slots, and delays in responding to market trends. Consumers are not disparate silos of preference based on the device they are using, but the market for advertising treats them that way due to limitations in the available technology tools, most of which are incapable of quickly and accurately integrating disparate data sets. For example, today, TV consumption data exists separately from set top box owner data and TV OEMs. It follows that advertising strategies for TV are planned

according to TV-specific criteria, and web and mobile advertising, which include sub categories such as social media, are each planned separately.

Furthermore, across the advertising industry, there are separate entities planning for different media platforms, such as set-top box, phone and desktop. Across the different media there exists disparate data, data systems, and data sources (vendors). Today, these device and media categories remain largely segmented when incorporated into advertising campaign strategies and planning.

Overall, then, the channels of commerce between an advertiser and the target audience are unduly complex: the route from generation of advertising content through to the ultimate consumer is labyrinthine, involving numerous players—some of whom are specific to certain media types, others of whom offer only a very specific and limited service in the supply chain. Advertisers would prefer to be able to access inventory with as few points of sale as possible. Correspondingly, the larger media organizations, such as cable companies, would benefit from an ability to access consumer data in such a way that they can make available to advertisers a range of tailored advertising inventory.

Currently, some companies attempt to link together a set of devices related to a particular consumer, but they are not capable of treating the disparate data sources with any reliable level of data integration, or at a scale that is useful to advertisers. Identifying selections of devices by comparing and modeling incomplete user data against that of other similar users within the market segment is a partial solution to this issue, but current methods are not able to create associations at a level of granularity that is reliable or useful.

Today, probabilistic and deterministic methods are not widely utilized to associate mobile and computer devices to a precise audience, or household. One reason such methods are not more widely adopted is because of inefficient processing and pairing of data across different devices. For example, in order to predict consumer purchasing, viewing, and advertising interaction habits at a 1:1 level of an association between a user and their respective device, it is insufficient to assume that any single instance of device access is representative of that user's purchasing intent. This is due to modern day habits of media consumption—users consume media on a large variety of devices as well as via different media (such as Hulu, Netflix, or cable television). As such, a much more complex analysis that enables insight into the intersection of media consumption and a user's family of devices is necessary.

Another reason probabilistic and deterministic methods are not more readily available to gauge consumer purchasing habits is because access to user device data is not easily achieved. For example, under consumer privacy laws, it is unlawful to access a user's device without their explicit consent. Therefore, on a mass scale, it is often unknown what combination of devices a demographic of users use, and what media they consumed on their respective devices. This poses a large challenge for advertisers when determining which advertising inventory to purchase and how best to reach their target audience efficiently on a given category of device.

Today, the data systems that track consumer information for use in advertising targeting lack the ability to broadly combine and integrate transmutable and non-transmutable categories (i.e., requiring the integration of a plurality of membrane levels) of consumer data. Most data systems contain static, one-dimensional, homogeneous classifications of consumers. For example, a 29 year old who bought a car two years ago will be a consumer data point that will

not adjust or be updated over time. While adjusting the age of this individual overtime time is simple, other transmutable characteristics such as desire to get married, pregnancy, or other lifestyle changes are not easy to assess or predict.

Accordingly, there is a need for a method of integrating and connecting data on a given consumer that is acquired over time from multiple different devices, and to use that integrated data in making reliable placement of advertising content across multiple devices. There would correspondingly be a benefit if that level of integration gave rise to a simpler transaction between advertiser and distributor.

The discussion of the background herein is included to explain the context of the technology. This is not to be taken as an admission that any of the material referred to was published, known, or part of the common general knowledge as at the priority date of any of the claims found appended hereto.

Throughout the description and claims of the application the word “comprise” and variations thereof, such as “comprising” and “comprises”, is not intended to exclude other additives, components, integers or steps.

#### SUMMARY

The instant disclosure addresses the processing of consumer data and advertising inventory in connection with optimizing placement of advertising content across display devices including one or more TVs. In particular, the disclosure comprises methods for doing the same, carried out by a computer or network of computers. The disclosure further comprises a computing apparatus for performing the methods, and computer readable media having instructions for the same. The apparatus and process of the present disclosure are particularly applicable to video content in online and TV media.

In overview, the present method allows an advertiser to allocate media strategies to different inventory types based on a probability of matching to an audience category or type. In particular, the present technology relates to systems and methods for optimizing advertising campaigns, as well as to scenario modeling as a way to improve yield. The methods herein lend themselves to increasing return on investment of money spent on advertising by increasing efficiency and reducing costs associated with identifying an advertising strategy, as well as improving yield for multichannel video programming distributors.

In an alternative embodiment, the system develops an advertising strategy designed around specified campaign (end-user) parameters. The strategy is directed to advertising occurring on TV, video-on-demand, display ads, and within mobile and desktop environments.

The method involves analyzing consumer, media and related data, from an unlimited number of data inputs, including but not limited to: behavioral such as specific viewing and purchasing histories of individual consumers, as well as demographic, and location-related sources.

The present technology includes programmatic generation of look-alike models in a consumer's device graph to predict future consumption behavior. The method integrates actual content consumption behavior with broadcast user segments, as well as assigning the devices used for consumption to a single consumer.

The present disclosure provides for a method for allocating delivery of video advertising content to a consumer on a TV, the method comprising: receiving a pricepoint and one or more campaign descriptions from an advertiser, wherein each of the campaign descriptions comprises a schedule for

delivery of an item of advertising content across one or more TVs accessed by a consumer, and a target audience, wherein the target audience is defined by one or more demographic factors; identifying one or more hard constraints associated with the one or more campaign descriptions; defining a pool of consumers based on a graph of consumer properties, wherein the graph contains information about two or more TV and mobile devices used by each consumer, demographic and online behavioral data on each consumer and similarities between pairs of consumers, and wherein the pool of consumers comprises consumers having at least a threshold similarity to a member of the target audience; receiving a list of inventory from one or more content providers, wherein the list of inventory comprises one or more slots for TV and online; identifying one or more advertising targets, wherein each of the one or more advertising targets comprises a sequence of slots consistent with one or more of the campaign descriptions and the one or more hard constraints, and has an overall cost consistent with the pricepoint; carrying out an optimization of allocation of the advertising content of the one or more campaign descriptions to the one or more advertising targets based on one or more soft constraints, thereby producing one or more solutions; communicating a list of one or more solutions to the advertiser, wherein a solution comprises a match of the campaign description to one or more slots within TV content identified as likely to be viewed by the pool of consumers; and delivering the item of advertising content to a consumer in the pool of consumers via a first media conduit on a television.

The present disclosure further provides for computer readable media, encoded with instructions for carrying out methods described herein and for processing by one or more suitably configured computer processors.

The present disclosure additionally includes a computing apparatus configured to execute instructions, such as stored on a computer readable medium, for carrying out methods described herein.

#### BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows, diagrammatically, relationships between parties that contribute to delivery of advertising content, such as advertisers, an advertising exchange, media conduits, and consumers.

FIG. 2 shows another set of relationships between parties that utilize the technology herein.

FIG. 3 shows a consumer graph.

FIG. 4 shows a node in a consumer graph.

FIGS. 5 and 6 show steps in creation of a consumer graph.

FIG. 7 show a flow chart of a method herein.

FIG. 8 shows a flow-chart of a process as described herein.

FIG. 9 shows an apparatus for performing a process as described herein.

FIGS. 10A-10D show aspects of an exemplary user interface.

Like reference symbols in the various drawings indicate like elements.

#### DETAILED DESCRIPTION

The instant technology is directed to computer-implemented methods that combine actual content consumption behavior, segments of the consumer population, and an assignment of the devices used for media consumption, in order to provide an advertiser with more targeted options for

advertising inventory purchase. The methods are particularly suited for practice by multichannel video programming distributors (MVPD's), wherein the MVPD has access to both the consumer data and an ability to run the methods at the request of an advertiser.

#### Advertising Functions

Relationships between entities in the business of purchase, delivery and consumption of advertising content are depicted in FIG. 1. As can be seen, the advertising ecosystem is complex, involves many different entities, and many different relationships. The methods and technology herein can be used to simplify many of the relationships of FIG. 1.

An advertiser **101** is a purchaser of advertising inventory **109**. An advertiser may be a corporation that exercises direct control over its advertising functions, or it may be an agency that manages advertising requirements of one or more clients, usually corporate entities. The advertiser intends to make advertising content **103** (also an "advertisement" herein) available to one or more, typically a population of, consumers **105**, on one or more devices **107** per consumer.

Devices include, for a given consumer, one or more of: TV's (including SmartTV's), mobile devices (cell phones, smartphones, media players, tablets, notebook computers, laptop computers, and wearables), desktop computers, networked photo frames, set top boxes, gaming consoles, streaming devices, and devices considered to function within the "Internet of Things" such as domestic appliances (fridges, etc.), and other networked in-home monitoring devices such as thermostats and alarm systems.

The advertising content **103** has typically been created by the advertiser **101** or a third party with whom the advertiser has contracted, and normally includes video, audio, and/or still images that seek to promote sales or consumer awareness of a particular product or service. Advertising content **103** is typically delivered to consumers via one or more intermediary parties, as further described herein.

Advertising content is typically of two different types: branding, and direct-response marketing. The timeframe is different for these two types. Branding promotes awareness; direct response marketing is designed to generate an immediate response. For example, an automobile manufacturer may put out direct response marketing material into the market place, and wants to measure responses by who went to a dealership or website after seeing an advertisement. The methods herein can be applied to both types of advertising content, but the measurement of effectiveness is different for the two types: for example, effectiveness of branding is measured by GRP's (further described elsewhere herein), and results of direct response marketing can be measured by, for example, website visits.

When delivered to a mobile device such as a phone or a tablet, advertising content **103** may additionally or alternatively take the form of a text/SMS message, an e-mail, or a notification such as an alert, a banner, or a badge. When delivered to a desktop computer or a laptop computer or a tablet, the advertising content **103** may display as a pop-up within an app or a browser window, or may be a video designed to be played while other requested video content is downloading or buffering.

Consumers **105** are viewers and potential viewers of the advertising content **103** and may have previously purchased the product or service that is being advertised, and may—advantageously to the advertiser—be learning of the product or service for the first time when they view the advertising content **103**.

Advertising inventory **109** (also inventory or available inventory, herein) comprises available slots, or time slots

117, for advertising across the several media interfaces, or conduits 111, through which consumers access information and advertising content. Such media interfaces include TV, radio, social media (for example, online networks, such as LinkedIn, Twitter, Facebook), digital bill boards, mobile apps, and the like. Media conduits 111 may generate their own content 113, or may be broadcasting content from one or more other content providers or publishers 115. For example, a cable company is a media conduit that delivers content from numerous TV channel producers and publishers of content. Media interfaces may also be referred to as content providers, generally, because they deliver media content 113 (TV programs, movies, etc.) to consumers 105. One aspect of the technology herein includes the ability to aggregate inventory 109 from more than one type of media interface or content provider. Media conduits 111 also deliver advertising content 103 that has been purchased for delivery at time slots 117, to consumers 105 for viewing on various devices 107. A publisher 115 is typically a content owner (e.g., BBC, ESPN).

A slot 117 is a time, typically expressed as a window of time (1 minute, 2 minutes, etc.) at a particular time of day (noon, 4:30 pm, etc., or a window such as 2-4 pm, or 9 pm-12 am), or during a specified broadcast such as a TV program, on a particular broadcast channel (such as a TV station, or a social media feed). An available slot is a slot in the inventory that an advertiser may purchase for the purpose of delivering advertising content. Typically it is available because another advertiser has not yet purchased it. As further described herein, a slot may additionally be defined by certain constraints such as whether a particular type of advertising content 103 can be delivered in a particular slot. For example, a sports equipment manufacturer may have purchased a particular slot, defined by a particular time of day on a particular channel, and may have also purchased the right to exclude other sports equipment manufacturers from purchasing slots on the same channel within a certain boundary—in time—of the first manufacturer's slot. In this context, a "hard constraint" is a legal or otherwise mandatory limitation on placing advertising in particular time slots or within specified media. A "soft constraint" refers to desired (non-mandatory) limitations on placing advertising in particular time slots within specified media. "Constraint satisfaction" refers to the process of finding a solution to a set of constraints that impose conditions that the variables must satisfy. The solution therefore is a set of values for the variables that satisfies all constraints.

Information is intended to mean, broadly, any content that a consumer can view, read, listen to, or any combination of the same, and which is made available on a screen such as a TV screen, computer screen, or display of a mobile device such as a tablet, smart-phone, or laptop/notebook computer, a wearable such as a smart-watch, fitness monitor, or an in-car or in-plane display screen. Information is provided by a media interface 111 such as a TV or radio station, a multi-channel video programming distributor (such as a cable TV provider, e.g., Comcast), or an online network such as Yahoo! or Facebook.

VOD refers to video on demand systems, which allow users to select and watch or listen to video or audio content when they choose to, rather than having to watch content at a scheduled broadcast time. Internet technology is often used to bring video on demand to televisions and personal computers. Television VOD systems can either stream content through a set-top box, a computer or other device, allowing viewing in real time, or download it to a device

such as a computer, digital video recorder (also called a personal video recorder) or portable media player for viewing at any time.

The communication between the advertisers and the media conduits can be managed by up to several entities, including: a demand-side provider (DSP) 123, an advertising exchange 119, and a supply-side provider 121. An advertising exchange 119 (also, exchange herein) is an environment in which advertisers can bid on available media inventory. The inventory may be digital such as via online delivery over the Internet, or via digital radio such as SiriusXM, or may be analog, such as via a TV channel such as ESPN, CNN, Fox, or BBC, or an FM/AM radio broadcast. An advertising exchange 119 typically specializes in certain kinds of content. For example, SpotX specializes in digital content, WideOrbit specializes in programmatic TV.

Supply-side provider (SSP) 121 is an intermediary that takes inventory 109 from a media conduit 111, and makes it available to a demand-side provider (DSP) 123, optionally via exchange 119, so that advertisers can purchase or bid on the inventory when deciding how to position advertising content 103. SSP's have sometimes been categorized as public or private depending on whether a media conduit is able to limit the identity and number of advertisers who have access to the inventory. In some situations, an SSP interacts directly with a DSP without the need for an advertising exchange; this is true if the functions of an advertising exchange that a purchaser of advertising content relies on are performed by one or both of the DSP and SSP. The technology herein is particularly suited for being implemented and being carried out by a suitably-configured DSP.

In one configuration, an advertising exchange 119 interfaces between a supply side provider (SSP) 121 and a demand side provider (DSP) 123. The interfacing role comprises receiving inventory 109 from one or more SSP's 121 and making it available to the DSP, then receiving bids 125 on that inventory from the DSP and providing those bids 125 to the SSP. Thus, a DSP makes it possible for an advertiser to bid on inventory provided by a particular SSP such as SPotX, or WideOrbit. In some configurations, the DSP takes on most or all of the role of an advertising exchange.

An advertising campaign (or campaign) is a plan, by an advertiser, to deliver advertising content to a particular population of consumers. A campaign will typically include a selection of advertising content (such as a particular advertisement or various forms of an advertisement, or a sequence of related advertisements intended to be viewed in a particular order), as well as a period of time for which the campaign is to run (such as 1 week, 1 month, 3 months). An advertiser typically transmits a campaign description 127 to an advertising exchange 119 or a DSP 121, and in return receives a list of the inventory 109 available. A campaign description 127 may comprise a single item of advertising content 103 and one or more categories of device 107 to target, or may comprise a schedule for sequential delivery of two or more items of advertising content 103 across one or more devices 107. A campaign description 127 may also comprise a description of a target audience, wherein the target audience is defined by one or more demographic factors selected from, but not limited to: age range, gender, income, and location.

The DSP 123 then provides an interface by which the advertiser 101 can align its campaign descriptions 127 against inventory 109 and purchase, or bid on, various slots 117 in the inventory. The DSP 123, or an exchange 119, may be able to provide more than one set of inventory that

matches a given campaign description **127**: each set of inventory that matches a given campaign description is referred to herein as an advertising target **129**. The advertiser **101** may select from among a list of advertising targets, the target or targets that it wishes to purchase. Once it has purchased a particular target, the SSP **121** is notified and delivery instructions **137** are sent to the various media conduits **111** so that the advertising content **103** can be delivered in the applicable time slots **117**, or during selected content **113**, to the relevant consumers.

A purchase of a given slot is not simply a straightforward sale at a given price, but is achieved via a bidding process. The DSP will place bids on a number of slots, and for each one, will have identified a bid price that is submitted to the SSP. For a winning bid, the SSP delivers the advertising content to the media conduit, and ultimately the consumer. Bids are generally higher for specific targeting than for blanket targeting.

The bidding process depends in part on the type of advertising content. TV content can be scheduled in advance, whereas for online content, the typical bid structure is 'last-in-time' bidding: the advert is delivered only if a particular consumer is seen online. In general, the methods herein are independent of bidding process, and are applicable to any of the bidding methods typically deployed, including real-time-bidding, as well as bidding that exploits details of programmatic TV data.

By serving a tag with a given online ad, by using a protocol such as VPAID (<https://en.wikipedia.org/wiki/Mixpo>) or VAST (video advert serving template), the tag collects data including whether a consumer clicked on, or viewed, the content. The tag typically contains a number of items of data relating to how a consumer interacted with the advertising content. The items of data can be returned to the SSP and/or the DSP in order to provide feedback on the circumstances of delivery of the advertisement. For example, the items of data can include a datum relating to whether a user clicked on a video online. Certain items of data correspond to events that are referred to in the industry as "beacon" events because of their salience to an advertiser: for example a beacon event can include the fact that a user stopped a video segment before it completed.

The process of generating advertising targets may also depend one or more campaign requirements. A campaign requirement, as used herein, refers to financial constraints such as a budget, and performance specifications such as a number of consumers to target, set by an advertiser or other purchaser of advertising inventory. Campaign requirement information is used along with campaign descriptions when purchasing or bidding on inventory.

DSP's **123** also provide advertisers **101** with data on consumers and devices, aggregated from various sources. This data helps an advertiser choose from the inventory, those time slots and media conduits that will best suit its goals.

Data used by DSP's may include census data **131**, or data on specific consumers and devices **133**. Census data **131** includes data on a population that can be used to optimize purchase of inventory. Census data **131** can therefore include demographic data such as age distribution, income variations, and marital status, among a population in a particular viewing region independent of what media interfaces the members of the population actually view. Census data **131** can be aggregated from a variety of sources, such as state and county records, and U.S. Census Bureau data.

A data management platform (DMP) **135** can provide other types of third party data **133** regarding consumers and

the devices they use to the DSP. Typically a DMP provides a data warehousing facilities with embedded functionality. DMPs download data and can perform a variety of analytical functions ranging from sorting, storing, processing, applying matching algorithms, and providing data outputs to purchasers and subscribers. Examples of DMP's include: Krux, Exelate, Nielsen, Lotame. The consumer and device data **133** that is delivered to a DSP from a third party provider may complement other consumer and device data **143** that is provided by the media conduits. Data on consumers and the devices they use that is relevant to an advertiser includes matters of viewing habits as well as specific behavioral data that can be retrieved directly from a media conduit. For example, as further discussed elsewhere herein, when a media conduit serves an advertisement to a consumer, the conduit can collect information on that user's manner of access to the advert. Due to the volume of data involved, after a relatively short period of time, such as 14 days, a media conduit may not be able to furnish any information on a particular consumer. In that instance, the DSP can get data on that user from a third party such as a DMP. Third parties can get data offline as well. As used herein, an offline event is one that happens independently of the Internet or a TV view: for example, it can include purchase of an item from a store and other types of location-based events that an advertiser can view as significant. Data can be shared between the entities herein (e.g., between a DMP and a DSP, and between DSP and SSP, and between media conduits and a SSP or advertising exchange) using any commonly accepted file formats for sharing and transfer of data: these formats include, but are not limited to: JSON, CSV, and Thrift, as well as any manner of text file appropriately formatted.

#### Role of Multichannel Video Programming Distributor

FIG. 2 shows an alternative set of relationships between various entities, in which a multichannel video programming distributor (MVPD) plays a principal role.

In FIG. 2, advertisers and agencies as applicable **101** interact directly with MVPD's **211**, by delivering campaign descriptions **127**, advertising content **103**, and by making bids **125** on inventory **109**. The MVPD's have made available to the advertisers their inventory **109** (which is typically TV content **113**, but may also include online data, from publishers **115**) in order that the advertisers can bid on it. The MVPD's have amassed a database of consumer and device data **143** (e.g., from TV viewing habits) as well as the tools to analyze it. As described elsewhere herein, the MVPD assists the advertiser to determine the most appropriate slots in the inventory and then ensures that the advertising content **103** is delivered to a target population of consumers.

#### Definitions

An impression refers to any instance in which an advertisement reaches a consumer. On a TV, it is assumed that if the TV is broadcasting the advertisement then an individual known to be the owner of, or a regular viewer of, that TV will have been exposed to the advertisement, and that display counts as an impression. If multiple persons are in the same household then the number of impressions may equal the number of persons who can view that TV. In the online environment, an impression occurs if a consumer is viewing, say, a web-page and the advertisement is displayed on that web-page such as in the form of a pop-up, or if the user has clicked on a link which causes the advertisement to run.

An audience segment is a list of consumers, de-identified from their personally identifiable information using cookie syncing or other methods, where the consumers belong to a type (income, gender, geographic location, etc.), or are associated with a behavior: purchases, TV viewership, site visits, etc.

Online, as used herein, means connected to the Internet or another computer network that permits devices to communicate with one another, such as an intranet. A device may be online if it is accessing the Internet or other network via a wifi connection or over the cellular data network, or by using a short-range communication protocol such as bluetooth. Thus online devices typically include, but are not limited to: computers, such as personal or desktop computers, workstations, laptops, notebooks, and tablets; e-books such as the Nook; mobile devices such as tablets (Apple iPad, Samsung Galaxy, etc.) and cellular phones or network-accessible music players, mobile hotspots; displays in items of transportation such as cars, buses, trains, shuttles, and airplanes; wearables such as watches, fitness monitors, virtual reality viewers such as Oculus, and devices that come within the general category of the “internet of things”, such as home appliances, including but not limited to smart TV’s, fridges, digital photo-frames, thermostats, and security systems.

Cookie syncing refers to a process that allows data exchange between DMP’s SSP’s and DSP’s, and more generally between publishers of content and advertisement buyers. A cookie is a file that a mobile device or desktop computer uses to retain and restore information about a particular user or device. The information in a cookie is typically protected so that only an entity that created the cookie can subsequently retrieve the information from it. Cookie syncing is a way in which one entity can obtain information about a consumer from the cookie created by another entity, without necessarily obtaining the exact identify of the consumer. Thus, given information about a particular consumer received from a media conduit, through cookie syncing it is possible to add further information about that consumer from a DMP.

For mobile devices, there is a device ID, unique to a particular device. For TV’s there is a hashed IP address. The device ID information may be used to link a group of devices to a particular consumer, as well as link a number of consumers, for example in a given household, to a particular device. A DSP may gather a store of data, built up over time, in conjunction with mobile device ID’s and TV addresses that augment ‘cookie’ data.

Cross-screen refers to distribution of media data, including advertising content, across multiple devices of a given consumer, such as a TV screen, computer screen, or display of a mobile device such as a tablet, smart-phone or laptop/notebook computer, a wearable such as a smart-watch or fitness monitor, or an in-car, or in-plane display screen, or a display on a networked domestic appliance such as a refrigerator.

Reach is the total number of different people exposed to an advertisement, at least once, during a given period.

In a cross-screen advertising or media campaign, the same consumer can be exposed to an advertisement multiple times, through different devices (such as TV, desktop or mobile) that the consumer uses. Deduplicated reach is the number of different people exposed to an advertisement irrespective of the device. For example, if a particular consumer has seen an advertisement on his/her TV, desktop and one or more mobile devices, that consumer only contributes 1 to the reach.

The incremental reach is the additional deduplicated reach for a campaign, over and above the reach achieved before starting the campaign, such as from a prior campaign. In one embodiment herein, a type of campaign can include a TV extension: in this circumstance, an advertiser has already run a campaign on TV, but is reaching a point of diminished returns. The advertiser wants to find ways to modify the campaign plan for a digital market, in order to increase the reach. In this way, a DSP may inherit a campaign that has already run its course on one or more media conduits.

In addition to TV programming content, and online content delivered to desktop computers and mobile devices, advertisements may be delivered within OTT content. OTT (which derives from the term “over the top”) refers to the delivery of audio, and video, over the Internet without the involvement of a MVPD in the control or distribution of the content. Thus, OTT content is anything not tied to particular box or device. For example, Netflix, or HBO-Go, deliver OTT content because a consumer doesn’t need a specific device to view the content. By contrast, MVPD content such as delivered to a cable or set top box box is controlled by a cable or satellite provider such as Comcast, AT&T or DirecTV, and is not described as OTT. OTT in particular refers to content that arrives from a third party, such as Sling TV, YuppTV, Amazon Instant Video, Mobibase, Dramatize, Presto, DramaFever, Crackle, HBO, Hulu, myTV, Netflix, Now TV, Qello, RPI TV, Viewster, WhereverTV, Crunchyroll or WWE Network, and is delivered to an end-user device, leaving the Internet service provider (ISP) with only the role of transporting IP packets.

Furthermore, an OTT device is any device that is connected to the internet and that can access a multitude of content. For example, Xbox, Roku, Tivo, Hulu (and other devices that can run on top of cable), a desktop computer, and a smart TV, are examples of OTT devices.

Gross rating point (GRP) refers to the size of an advertising campaign according to schedule and media conduits involved, and is given by the number of impressions per member of the target audience, expressed as a percentage (GRP can therefore be a number >100. For example, if an advert reaches 30% of the population of L.A. 4 times, the GRP is 120. (The data may be measured by, e.g., a Nielsen panel of say 1,000 viewers in L.A.).

The target rating point (TRP) refers to the number of impressions per target audience member, based on a sample population. This number relates to individuals: e.g., within L.A. the advertiser wants to target males, 25 and older. If there are 100 such persons in the L.A. panel and 70% saw the ad., then the TRP is 70%×number of views.

“Cross-Screen” refers to analysis of media, consumer, and device data that combines viewer data across multiple devices.

“High frequency” refers to high frequency trading related to advertising purchases and sales. The methods and technology herein can be practiced by trading platforms for advertising that utilize computers to transact a large number of bid requests for advertising inventory at high speeds. The systems herein can operate at high-frequency, such as running from 10,000 to 100,000 queries of audience impressions per second. The queries can be dynamic, and in real-time.

Consumer Data

Data about consumers can be categorized into two groups: there are non-transmutable characteristics such as ethnicity, and gender; and there are transmutable characteristics such as age, profession, address, marital status, income, taste and preferences. Various transmutable characteristics such as



profession are subject to change at any time, while others such as age change at a consistence rate. Today, the data systems that track consumer information for use in targeting advertising content lack the ability to broadly track both categories of consumer data. Most data systems contain static, homogenous classifications of consumers. For example, a 29-year old who bought a car two years ago will be a consumer data point that will not be updated or augmented with time. Even if the age of the individual as stored in a system can be adjusted with time, other transmutable characteristics such as change in marital state, or lifestyle changes, are not taken into account in this consumer's classification.

At various stages of the methods herein, it is described that each consumer in a population of consumers is treated in a particular way by the method: for example, a computer may be programmed to analyze data on each consumer in its database in order to ascertain which, if any, have viewed a particular TV show, or visited a particular website; alternatively, some comparative analysis may be performed, in which attributes of each user in one category of population is compared with attributes of each consumer in another category of population. Each population set may comprise many thousands of individuals, or many hundreds of thousands, or even millions or many millions of individuals. It is assumed herein that the methods, when deployed on suitable computing resources, are capable of carrying out stated calculations and manipulations on each and every member of the populations in question. However, it is also consistent with the methods herein that "each consumer" in a population may also mean most consumers in the population, or all consumers in the population for whom the stated calculation is feasible. For example, where one or more given consumers in a population is omitted from a particular calculation because there is insufficient data on the individual, that does not mean that an insufficient number of members of the population is analyzed in order to provide a meaningful outcome of the calculation. Thus "each" when referencing a population of potentially millions of consumers does not necessarily mean exactly every member of the population but may mean a large and practically reasonable number of members of the population, which for the purposes of a given calculation is sufficient to produce a result.

#### Consumer Graph

A consumer graph is a graph in which each node represents a consumer (or individual user). The technology utilizes various implementations of a weighted graph representation in which relationships between consumers (nodes) are defined as degrees of similarity (edges). A consumer graph is used herein to categorize, store, and aggregate large amounts of consumer data, and allow an entity such as a DSP to make connections between data used to build a consumer graph with other data—such as TV viewing data—via data on given consumers' devices.

One way to construct the graph is by using deterministic relationship data; another is probabilistically using the attributes of each node. In some instances, a combination of deterministic and probabilistic methods can be used. In a deterministic, approach, which is relatively straightforward, the basis is having exact data on a consumer, such as login information from a publisher. Thus, if a person has logged in multiple times on different devices with the same ID, then it is possible to be sure that the person's identity is matched. However, such exact information may not always be available. By contrast, in a probabilistic approach, it is necessary to draw inferences: for example, if the same device is seen in the same location, or similar behavior can be attributed to

a given device at different times, then it possible to conclude that the device belongs to the same user.

In some embodiments herein, machine learning methods, and Bayesian and regression algorithms, are used to explore commonalities between consumers. Such methods are useful in situations where there is a finite number of parameters to be considered. In some other embodiments, techniques of deep learning are more useful in finding consumer similarities and constructing a consumer graph. Machine learning is a preferred technique for matching exact pieces of information, for example whether the same websites have been visited by two consumers, but deep learning can explore the details of a particular video or TV program—for example, by analyzing natural scene statistics—and thereby ascertain, for example, whether two adverts that were viewed by a given consumer have something in common beyond their subject matter. For example, two adverts may include the same actor and be liked by a consumer for that reason, even though the products portrayed have little in common.

In preferred embodiments, the device graph herein is based on probabilistic data. The probabilistic approach to graph construction uses behavioral data such as viewership habits to match up users.

In some embodiments, an entity such as a DSP, can construct a device graph; in other embodiments it can obtain, such as purchase, it from a another entity such as a DMP.

In various embodiments herein, both a device graph and a consumer graph are operating together in a manner that permits tying in mobile data to TV data.

The term graph is used herein in its mathematical sense, as a set  $G(N, E)$  of nodes ( $N$ ) and edges ( $E$ ) connecting pairs of nodes. Graph  $G$  is a representation of the relationships between the nodes: two nodes that are connected by an edge are similar to one another according to some criterion, and the weight of an edge defines the strength of the similarity. Pairs of nodes that do not meet the similarity criterion are not joined by an edge. FIG. 3 illustrates graph concepts, showing 6 nodes,  $N_i$ — $N_e$ , in which three pairs of nodes are connected by edges.

In the implementation of a graph herein, a node,  $N$ , is an entity or object with a collection of attributes,  $A$ . In FIG. 3, each node has associated with it an array of attributes, denoted  $A_i$  for node  $N_i$ .

In the implementation of a graph herein, an edge,  $E$ , existing between two nodes indicates the existence of a relationship, or level of similarity, between the two nodes that is above a defined threshold. The weight of an edge,  $w_E$ , is the degree of similarity of the two nodes. The weights of the edges in FIG. 3 are shown diagrammatically as thicknesses (in which case,  $w_{E12} > w_{E34} > w_{E15}$ ).

In a consumer graph, a node represents an individual, or a household comprising two or more individuals, with a set of attributes such as the gender(s) and age(s) of the individual(s), history of TV programs watched, web-sites visited, etc.

FIG. 4 illustrates an exemplary structure of a node of a consumer graph. Each node has a collection of attributes that include types and behaviors, for which data is continuously collected from first party and third party sources. Many of the attributes are transmutable if new information for the consumer becomes available, and the collection of attributes (i.e., the number of different attributes stored for a given consumer) can also grow over time as new data is collected about the consumer. An aspect of the technology herein is that the graph is constructed from a potentially unlimited number of inputs for a given consumer, such as online, offline, behavioral, and demographic data. Those inputs are

updated over time and allow the data for a given consumer to be refined, as well as allow the population of consumers on which data can be used to be expanded. The fact that there is no limit to the type and character of data that can be employed means that the methods herein are superior to those employed by panel companies, which rely on static datasets and fixed populations.

Some of the sources from which data is collected are as follows.

Type data is categorical data about a consumer that normally does not change, i.e., is immutable. Behavioral data is continuously updated based on a consumer's recent activity.

Each node includes a grouping of one or more devices (desktop, mobile, tablets, smart TV). For each device, data on the type of the user based on the device is collected from third party and first party sources.

Table 1 shows examples of data by category and source.

TABLE 1

1 <sup>st</sup> Party		3rd Party
Non-transmutable		Census (Govt.) Household income Education level (e.g., from Exelate) Gender (e.g., from Nielsen, DAR)
Transmutable	Behavior (online) TV viewing Viewability (e.g., how much of advert seen, kept on, visible online?) Online sites visited Location events	Offline behavior Retail Purchases Offsite visits (visited pharmacy, movie theater, car dealership, etc.)

First party data comprises data on a user's behavior, for example: purchases, viewership, site visits, etc., as well as types such as income, gender, provided directly by a publisher to improve targeting and reporting on their own campaigns. (For example, the Coca Cola company might provide to a DSP, a list of users who "like" Coke products on social media to improve their video advertising campaigns.) First party type data can be collected from advertisements served directly to the device, and from information collected from the device, such as one or more IP addresses. First party type data includes location from IP address, geolocation from mobile devices, and whether the device is located in a commercial or residential property.

Third party type data is obtained from external vendors. Through a one-on-one cookie synchronization or a device synchronization, an external vendor, for example a DMP such as Krux (<http://www.krux.com/>), Experian (which provides purchase behavior data), or Adobe, provides information about the cookie or device. Example data includes market segment occupied by the consumer, such as age range, gender, income level, education level, political affiliation, and preferences such as which brands the consumer likes or follows on social media. Additionally, external vendors can provide type data based on recent purchases attributed to the device. Third party data includes information such as gender and income because it is not collected directly from external vendors. Third party data can be collected without serving an advertisement. TV programs viewed and purchases are third party data.

First Party data is typically generated by a DSP; for example, it is data that the DSP can collect from serving an

Ad or a Brand/Agency that provides the data. First party data includes data that depends on having served an Ad to have access to it.

Behavioral data can be collected from the devices through first party and third party sources. Behaviors are first party data typically, and are mutable.

First party behavioral data is collected from advertisements served directly to the device. This includes websites visited, and the TV program, or OTT, or video on demand (VOD) content viewed by the device.

Third party behavioral data is obtained from external vendors, typically DMP's such as Experian, Krux, Adobe, Nielsen and Comscore, and advertising exchanges or networks, such as Brightroll, SpotX, FreeWheel, Hulu. Example data includes the history of TV programming viewed on the device in the last month, the history of websites visited by a personal computer or laptop, or mobile device, and history of location based events from mobile devices (for example, whether the device was at a Starbucks). In some instances, the same types of data can be obtained from both first party and third party entities.

Edges between the nodes in the consumer graph signify that the consumers have a threshold similarity, or interact with each other. The edges can be calculated deterministically, for example, if the nodes are in physical proximity, or probabilistically based on similarity in attributes. Probabilistic methods utilized include, but are not limited to: K-means clustering, and connected components analysis (which is based on graph traversal methods involving constructing a path across the graph, from one vertex to another. Since the attributes are transmutable, the edges can also change, either in their weighting or by being created or abolished if the similarity score for a pair of nodes alters. Thus the graph is not static, and can change over time. In some embodiments, change is dynamic: similarity scores are continually recalculated as nodes attributes for nodes are updated.

Typically, attributes and data are added dynamically (as they are obtained). The graph may be re-constructed weekly to take account of the new attributes and data, thereby establishing new weightings for the edges, and identifying newly connected or reconnected devices. (Graph construction and reconstruction may be done in the cloud, i.e., by distributing the calculations over many processors on a computer network, or on processors warehoused at a data-center under the control of the DSP.)

The similarity,  $S$ , between two nodes  $N_1$ ,  $N_2$ , is calculated according to a similarity metric, which is the inverse of a distance function,  $f(N_1, N_2)$ :  $N_1, N_2 \rightarrow S$ , that defines the similarity of two nodes based on their attributes.

In a consumer graph, similarity represents the likeness of two individuals in terms of their demographic attributes and their viewing preferences. Similarities can be calculated, attribute by attribute, and then the individual similarity attributes weighted and combined together to produce an overall similarity score for a pair of nodes.

When the attributes of two nodes are represented by binary vectors, there are a number of metrics that can be used to define a similarity between a pair of nodes based on that attribute. Any one of these metrics is suitable for use with the technology herein. In some embodiments, for efficiency of storage, a binary vector can be represented as a bit-string, or an array of bit-strings.

When working with a similarity metric that is the inverse of a distance function,  $f(N_i, N_j)$ , a zero value of the distance function signifies that the types and behaviors of the two nodes are identical. Conversely, a large value of the

distance function signifies that the two nodes are dissimilar. An example of a distance function is Euclidean distance,

$$f(N_i, N_j) = \|A_i - A_j\|^2$$

where  $A_i$ , and  $A_j$  are the sparse vectors representing the attributes of nodes  $N_i$  and  $N_j$ , and the distance is computed as a sum of the squares of the differences of in the values of corresponding components of each vector.

Comparisons of binary vectors or bit-strings can be accomplished according to one or more of several similarity metrics, of which the most popular is the Tanimoto coefficient. Other popular metrics include, but are not limited to: Cosine, Dice, Euclidean, Manhattan, city block, Euclidean, Hamming, and Tversky. Another distance metric that can be used is the LDA (latent Dirichlet allocation). Another way of defining a distance comparison is via a deep learning embedding, in which it is possible to learn the best form of the distance metric instead of fixing it as, e.g., the cosine distance. An example approach is via manifold learning.

The cosine dot product is a preferred metric that can be used to define a similarity between the two nodes in a consumer graph. The cosine similarity, that is the dot product of  $A_i$  and  $A_j$ , is given by:

$$f(N_i, N_j) = A_i \cdot A_j$$

In this instance, the vectors are each normalized so that their magnitudes are 1.0. A value of 1.0 for the cosine similarity metric indicates two nodes that are identical. Conversely, the nearer to 0.0 is the value of the cosine metric, the more dissimilar are the two nodes. The cosine metric can be converted into a distance-like quantity by subtracting its value from 1.0:

$$f(N_i, N_j) = 1 - A_i \cdot A_j$$

An example of a more complex distance function is a parameterized Kernel, such as a radial basis function.

$$f(N_i, N_j) = \exp(-\|A_i - A_j\|^2 / s^2),$$

where  $s$  is a parameter.

In the more general case in which the bit-string is a vector that contains numbers other than 1 and 0 (for example it contains percentages or non-normalized data), then one can calculate similarity based on distance metrics between vectors of numbers. Other metrics, such as the Mahalanobis distance, may then be applicable.

Typically, a similarity score,  $S$ , is a number between 0 and 100, though other normalization schemes could be used, such as a number between 0 and 1.0, a number between 0 and 10, or a number between 0 and 1,000. It is also possible that a scoring system could be un-normalized, and simply be expressed as a number proportional to the calculated similarity between two consumers.

In some embodiments, when calculating a similarity score, each contributing factor can be weighted by a coefficient that expresses the relative importance of the factor. For example, a person's gender can be given a higher weighting than whether they watched a particular TV show. The weightings can be initially set by application of heuristics, and can ultimately be derived from a statistical analysis of advertising campaign efficacy that is continually updated over time. Other methods of deriving a weighting coefficient used to determine the contribution of a particular attribute to the similarity score include: regression, or feature selection such as least absolute shrinkage and selection operator ("LASSO"). Alternatively, it is possible to fit to "ground truth data", e.g., login data. In some embodiments, as the system tries different combinations or features, which

one leads to greater precision/recall can be deduced by using a "held out" test data set (where that feature is not used in construction of the graph).

Another way of deriving a similarity score for a feature is to analyze data from a successive comparison of advertising campaigns to consumer feedback using a method selected from: machine learning; neural networks and other multi-layer perceptrons; support vector machines; principal components analysis; Bayesian classifiers; Fisher Discriminants; Linear Discriminants; Maximum Likelihood Estimation; Least squares estimation; Logistic Regressions; Gaussian Mixture Models; Genetic Algorithms; Simulated Annealing; Decision Trees; Projective Likelihood; k-Nearest Neighbor; Function Discriminant Analysis; Predictive Learning via Rule Ensembles; Natural Language Processing, State Machines; Rule Systems; Probabilistic Models; Expectation-Maximization; and Hidden and maximum entropy Markov models. Each of these methods can assess the relevance of a given attribute of a consumer for purposes of suitability for measuring effectiveness of an advertising campaign, and provide a quantitative weighting of each. Representation

To properly assess an entire population of consumers, a large number of nodes needs to be stored. Additionally, the collection of attributes that represent a node's types and behaviors can be sizeable. Storing the collection of the large number of attributes for the nodes is challenging, since the number of nodes can be as many as hundreds of millions. Storing the data efficiently is also important since the graph computations can be done most quickly and efficiently if the node data is stored in memory.

In a preferred embodiment, attributes are represented by sparse vectors. In order to accomplish such a representation, the union of all possible node attributes for a given type is stored in a dictionary. Then the type, or behavior, for each node is represented as a binary sparse vector, where 1 and 0 represent the presence and absence of an attribute, respectively. Since the number of possible attributes of a given type is very large, most of the entries will be 0 for a given consumer. Thus it is only necessary to store the addresses of those attributes that are non zero, and each sparse vector can be stored efficiently, typically in less than  $1/100$ th of the space that would be occupied by the full vector.

As an example, let the attributes encode the TV programs that a given consumer has viewed in the last month. The system enumerates all possible TV shows in the dictionary, which can be up to 100,000 different shows. For each node, whether the consumer watched the show in the last month is indicated with a 1, and a 0 otherwise.

If the attributes indicate different income levels, multiple income levels are enumerated, and a 1 represents that the consumer belongs to a particular income level (and all other entries are 0).

Thus for a consumer,  $i$ , having an annual income in the range \$30,000-\$60,000, and who has viewed "Top Gear" in the last month, the following is established:

TV\_Dictionary={"Walking Dead", "Game of Thrones", . . . "Top Gear"}

TV<sub>*i*</sub>=[0,0, . . . ,1]

TV<sub>*i*</sub> can be stored as simply [4]; only the 4th element of the vector is non-zero. Similarly, for income:

Income\_Dictionary={<\$30,000,\$30,000-\$60,000,\$60,000-\$100,000,>\$100,000}

Income<sub>*i*</sub>=[0,1,0,0]

Income<sub>i</sub> can be stored as simply [2], as only the second element of the vector is non-zero.

All the attributes of a node, *i*, can thus be efficiently represented with sparse vectors. This requires 2 to 3 orders of magnitude less memory than a dense representation.

#### Graph Construction

FIGS. 5 and 6 illustrate a flow-chart for steps in construction of a consumer graph. A MVPD can construct a graph using their data exclusively or they can take in a consumer graph from a third party, such as a DSP, or they can base a graph on a combination of data from both sources.

Initially, the graph is a collection of devices, which are mapped to consumers. Multiple data sources are used to group multiple devices (tablet, mobile, TV, etc.) to a single consumer. This typically utilizes agglomerative techniques. In order to attribute a single device (e.g., a Smart TV) to multiple consumers, a refinement technique is used.

With agglomerative methods, multiple devices can be grouped to a single consumer (or graph node). Some data sources used for this include, but are not limited to:

IP addresses: multiple devices belonging to same IP address indicates a single consumer or a household.

Geolocation: multiple devices that are nearby, using latitude and longitude, can be attributed to a single consumer.

Publisher logins: if the same consumer is logged in from multiple devices, those devices can be associated with that consumer.

During this process, the consumer's identity is masked, to obviate privacy concerns. The result is a single consumer ID that links particular devices together.

Let  $P(d_i, d_j)$  be the probability that the two devices,  $d_i$  and  $d_j$ , belong to the same node (consumer, or household). From multiple datasets obtained from different categories of device, it is possible to construct the probability:

$$P(d_i, d_j) = \frac{w_{IP} \times P(d_i, d_j | IP) \times w_{Geo} \times P(d_i, d_j | Geo) \times w_{Login} \times P(d_i, d_j | Login)}{Z}$$

where "×" means "multiply", where  $w_{\_}$  are weighting factors,  $P(d_i, d_j | Y)$  is a conditional probability (the probability of observing device *i* and device *j* belong to same user, if *Y* has the same value for both, and *Z* is a normalizing factor. Thus, *Y* may be an IP address. (The value of the conditional probability may be 0.80). Each data source gets a different weighing factor: for example, login data can be weighted higher than IP addresses. The weights can be fixed, or learned from an independent validation dataset.

Once multiple devices are grouped to a single node, the Types and Behaviors from the respective devices are aggregated to the singular node's attributes. For example, attributes (and the corresponding sparse vectors) from mobile (such as location events), and desktop (recent purchases) are aggregated. This provides more comprehensive information for a consumer, permitting more accurate and meaningful inferences for a node to be made.

Associating a device with a given consumer is possible due to the data that is associated with those devices and known to various media conduits. For examples, a Smart-TV stores location information as well as subscription information about the content broadcast by it. This information is shared with, and can be obtained from, other entities such as a cable company. Similarly, a mobile device such as a tablet or smartphone may be associated with the same (in-home) wifi network as the Smart-TV. Information about the location is therefore shared with, e.g., the cell-phone carrier, as well as broadcasters of subscription content to the mobile device. A key aspect of the graph methodology

herein is that it permits consumer information to be linked across different device and media platforms that have typically been segregated from one another: in particular, the graph herein is able to link consumer data from online and offline purchasing and viewing sources with TV viewing data.

With refinement methods, a single device (for example, a smart TV) can be associated with multiple consumers (or graph nodes) who, for example, own mobile devices that are connected to the same wifi network as the smart-TV.

Given a node, *n*, to which are assigned multiple devices, the various attributes are clustered into smaller groups of devices, for example, a TV ID, connected to multiple devices from a common IP address. The TV viewership data is aggregated along with the attributes from all the devices. A clustering algorithm, such as k-means clustering, can be applied to group the devices into smaller clusters. The number of clusters, *k*, can be set generally by the number of devices (by default  $k = \# \text{ number of devices} / 4$ ). Sometimes it is possible to only collect aggregate data at a household level. For example, there may be as many as 20 devices in one household. But by using behavioral data, it can be ascertained that the 20 devices have 4 major clusters, say with 5 devices each, where the clusters correspond to different individuals within the same household. Thus, although there are two categories of device (shared and personal), it is still important to attribute behavioral data to users.

Once a shared device is attributed to multiple nodes, the data collected from the device can be attributed to the nodes. For example, TV viewing data from a Smart TV can be collected from the OEM. Through this attribution, the TV viewing data can be added to the collection of a node's attributes. Ultimately, a Smart-TV can be attributed to different persons in the same household.

#### Lookalike Modeling by Learning Distance Functions

Given a graph,  $G(N, E)$ , and a functional form that defines a similarity metric, and a set of seed nodes, it is possible to generate a set of "lookalike" nodes that are similar to the seed nodes, where similarity is defined by a function that is fixed, or learned. This is useful when identifying new consumers who may be interested in the same or similar content as a group of consumers already known to an advertiser. Similar principles can be utilized when projecting likely viewing behavior of consumers from historical data on a population of consumers.

Seed nodes can be a set of nodes, e.g., household(s) or individual(s), from which to generate a set of lookalike nodes using a fixed, or learned, similarity metric. For example, seed nodes can be defined as an audience segment (such as list of users that saw a specific show for certain). This is useful for determining, for each member of the audience segment, a list of other audience members who might have similar viewing habits even if they did not watch exactly the same show as the seeds.

Given the set of seed nodes in a graph (and their attributes), the output of lookalike modeling is a set of nodes (that includes the seed nodes) that are similar to the seed nodes based on the fixed or learned similarity metric.

Several different vectors can be used in determining look-alike models: One is the vector of TV programs in total. This vector can be as long as 40 k elements. Another vector is the list of consumers who saw a particular program (e.g., The Simpsons). The vector of viewers for a given TV program can be as long as 10 M elements, because it contains one element per consumer. Another vector would be a vector of web-sites visited (say 100 k elements long).

Still another vector would be based on online videos viewed (which can also be 100 k elements long).

In general, TV program comparison data accesses a 10 M user base. Online data can identify a potentially much larger audience, such as 150 M consumers. It should be understood that TV data can be accumulated across a variety of TV consumption devices that include, but are not limited to linear, time-shifted, traditional and programmatic.

The similarity between 2 distinct nodes can be calculated from their attributes, represented by sparse vectors. Given a distance function  $f(N_i, N_j)$ , and a set of seed nodes,  $N_S$ , the pairwise distances between each element of the seed nodes,  $n$  in  $N_S$ , and all other nodes other than the seed node,  $n'$ , are calculated. That is, all quantities  $f(n, n')$  are calculated.

After calculating all pairwise similarities, only the nodes such that  $f(n, n') < T$  are selected.  $T$  is a threshold maximum distance below which the nodes are deemed to be similar. Alternatively, values of  $f(n, n')$  (where  $n$  is not  $n'$ ) are ranked in decreasing order, and the top  $t$  node pairs are selected. In either case,  $T$  and  $t$  are parameters that are preset (provided to the method), or learned from ground truth or validation data. The set of all nodes  $n'$  that satisfy the criteria above, form the set of "lookalike nodes".

#### Graph Inference

Given a graph  $G(N, E)$ , it is also possible to infer likely attributes of a node,  $n$ , based on the attributes of its neighbors in the graph. This can be useful when incomplete information exists for a given consumer but where enough exists from which inferences can be drawn. For example, TV viewership attributes may be missing for a node  $n$  (in general, there is either positive information if a user did watch a show, or it is unknown whether they watched it), whereas those attributes are available for neighbor nodes  $n'$ ,  $n''$  in the graph. Nodes  $n$ ,  $n'$ , and  $n''$  contain all other attributes, such as income level and websites visited.

In another example, it can be useful to calculate the probability that the consumer associated with node  $n$  would watch the show "Walking Dead", given that  $n'$ ,  $n''$  both also watch "Walking Dead". If the similarity, given by the weight of the edges between  $n$  and  $n'$ ,  $n''$ , are  $w'$ ,  $w''=0.8$  and  $0.9$  respectively, and the likelihood of  $n$  watching the show based on his/her own attributes is  $0.9$ , then the probability is given by:

$$P(n \text{ watches "Walking Dead"}) = [0.8 \times 0.9 + 0.9 \times 0.9] / [0.8 \times 0.9 + 0.9 \times 0.9 - F(1 - 0.8 \times 0.9) - F(1 - 0.9 \times 0.9)] = 0.765$$

Similar principles can be utilized when projecting likely viewing behavior of consumers from historical data on a population of consumers.

#### Accuracy

The graph is continually refined as new data is received. In one embodiment, a technique such as machine learning is used to improve the quality of graph over time. This may be done at periodic intervals, for example at a weekly build stage. It is consistent with the methods herein that the graph utilized is updated frequently as new consumer data becomes available.

To determine the accuracy of the graph, the precision and recall can be compared against a validation dataset. The validation dataset is typically a (sub)graph where the device and node relationships are known with certainty. For example, the login information from an online network such

as eHarmony, indicates when the same user has logged into the site from different desktops (office, laptop), and mobile devices (smartphone and tablet). All the devices that are frequently used to login to the site are thus tied to the same consumer and thereby that individual's graph node. This information can be used to validate whether the constructed graph ties those devices to the same node.

If  $D$  is the set of devices in the validation set, let  $Z(D)$  denote the graph, consisting of a set of nodes, constructed from the set of devices,  $D$ . For different datasets, and different graph construction methods, it is possible to obtain different results for  $Z(D)$ .

For the set  $Z(D)$ , true positive (TP), false positive (FP), and false negative (FN) rates can all be calculated. True positives are all nodes in  $Z(D)$  that are also nodes in the validation set. False positives are all nodes in  $N(D)$  that do not belong to the set of nodes in the validation set. False negatives are all nodes that belong to the validation set, but do not belong to  $Z(D)$ .

Precision, defined as  $TP/(TP+FP)$ , is the fraction of retrieved devices that are correctly grouped as consumer nodes.

Recall, defined as  $TP/(TP+FN)$ , is the fraction of the consumer nodes that are correctly grouped.

Depending on the application at hand, there are different tradeoffs between precision and recall. In the case of constructing a consumer graph, it is preferable to obtain both high precision and high recall rates that can be used to compare different consumer graphs.

The validation dataset must not have been used in the construction of the graph itself because, by doing so, bias is introduced into the precision and recall values.

#### Learning the Similarity Metric

Another feature of the graph that can be adjusted as more data is introduced is the underlying similarity metric. Typically, the metric is fixed for long periods of time, say 5-10 iterations of the graph, and the metric is not reassessed at the same frequency as the accuracy.

In the case where the distance function is not fixed, it is possible to learn the parameters of a particular distance function, or to choose the best distance function from a family of such functions. In order to learn the distance function or its parameters, the values of precision and recall are compared against a validation set.

Suppose a goal is to predict the lookalike audience segment that are high income earners, based on the attributes of a seed set of known high income earners. The similarity of the seed nodes to all other nodes in the graph is calculated for different distance functions, or parameters of a particular distance function. The distance function uses the attributes of the nodes, such as online and TV viewership, to calculate the similarities.

For example, if the distance function is the radial basis function with parameter,  $s$ :

$$f(N_i, N_j) = \exp(-\|A_i - A_j\|^2 / s^2),$$

then the pairwise distances from the seed nodes to all other nodes, are calculated for different values of  $s$ , using the same threshold distance value,  $T$ , to generate the set of lookalike nodes. For different values of  $s$  (the parameter that needs to be learned), the calculations produce different sets of lookalike nodes, denoted by  $N_S(s)$ .

For the set  $N_S(s)$ , it is possible to calculate true positive (TP), false positive (FP) and false negative (FN) rates. True positives are all nodes in  $N_S(s)$  that also belong to the target set in the validation set. In this example, all the nodes that are also high income earners (in ground truth set). False

positives are all nodes in  $N\_S(s)$  that do not belong to the target set (not high income earners). False positives are all nodes in  $N\_S(s)$  that do not belong to the target set (not high income earners). False negatives are all nodes that belong to the validation set (are high income earners), but do not belong to  $N\_S(s)$ .

Based on the application, it is possible to require different tradeoffs between precision and recall. In the case of targeting an audience with an advertisement, a high recall rate is desired, since the cost of exposure (an advertisement) is low, whereas the cost of missing a member of a targeted audience is high.

In the example herein, the aim is to choose the value of  $s$  for which both the precision and recall rates are high from amongst possible values of  $s$ . For other types of distance function, there may be other parameters for which to try to maximize the precision and recall rates.

The accuracy of a lookalike model can only be defined for a target audience segment. For example, it is possible to predict whether a lookalike segment also comprises high income earners, from a seed set of high income earners using TV viewing and online behavior datasets. Predictions can be validated using a true set of income levels for the predicted set of nodes. This gives the accuracy of the predictions. However, the accuracy of predictions for one segment are not meaningful for a new target segment, such as whether those same users are also luxury car drivers.

#### Calculating Deduplicated Reach

The consumer graph connects a node (consumer) to all the devices that he or she uses. Thus the graph enables deduplicating the total exposure to an advertisement, to individuals. For example, if user abc123 has already seen a particular advertisement on each of his TV, desktop and mobile device, the total deduplicated exposures will count as 1. This enables the calculation of the following metrics for direct measurement.

The deduplicated exposed audience is the number of users belonging to the target audience segment in the consumer graph who were exposed to the advertisement after deduplication. Then, the direct deduplicated reach is:

$$\text{Deduplicated Reach} = \frac{\text{Deduplicated Exposed Audience}}{\text{Total Audience}}$$

For sampled measurement, this enables the calculation of the deduplicated exposed sampled audience as the number of sampled users who belong to the target audience segment who were exposed to the advertisement after deduplication. Then, the sampled reach is:

$$\text{Deduplicated Sampled Reach} = \frac{\text{Deduplicated Exposed Sampled Audience}}{\text{Total Sampled Audience}}$$

In the case of modeled measurement data, the ID of the user in the consumer graph from whom the data was collected is not known. Hence, the reach data cannot be deduplicated on a one-to-one level.

Calculation of deduplicated reach can be useful in managing targeting, if an advertiser wants to impose a frequency cap on consumers (for example, if the advertiser doesn't want to show the same advert to the same user more than twice). Deduplicated reach also provides a convenient metric by which to optimize the efficacy of an advertising campaign: for example, by calculating the deduplicated reach over time, as an advertising campaign is adjusted, improvements can continue to be made by altering parameters of the campaign such as, for example, consumer demographic, or time and channel of broadcast of TV content.

#### Calculating Incremental Reach

On day  $t$ , let the deduplicated reach (direct or sampled) be  $x$ . The incremental reach is the additional deduplicated reach after running the campaign. In a cross-screen environment, this is a useful parameter to calculate if an advertiser wants to be able to assess whether they can extend a 30% reach via TV to say, a 35% reach by extending to mobile platforms. One caveat is that in direct measurement of, e.g., TV data, the portion of the sample obtained for smart-TV's is only a subset of the overall data, due to the relatively small number of smart-TV's currently in the population at large.

In the case of modeled measurement data such as is obtained from a panel where the nature of the sample has to be inferred, the ID of the user in the consumer graph from whom the data was collected is not known. Hence, it is not possible to tell if the same user has viewed the advertisement in the past. Therefore the incremental deduplicated reach cannot be calculated for modeled data because devices cannot be associated with particular users. Since the incremental reach from the sampled measurement, without deduplication, can be calculated, as described above, the methods herein are superior to panel-based methods.

#### Modeling of User Device Habits

Solving the issue of disparate data tracking requirements involves processing data inputs that may have sequencing tags that reflect the nature of each category of device or medium. In one embodiment, the technology herein addresses this issue by allocating data based on the nature of each device. For example, a first batch of consumer data is limited to data mapped from the specific devices that a specific consumer uses, the sources being a plurality of third party APIs. This device data updates such aspects as when, where, how, how long, and the level of engagement by the consumer, on each particular device. The data is then integrated and processed separately from other consumer data. All device usage data is integrated to create a more precise understanding of the access points and behavior of the consumer over time.

In another example, a batch of consumer data is based on consumption data. Within the data stored for each device, exists the potential to include a further plurality of third party data concerning actual media consumed. This data can be obtained from content providers, OEMs, publishers, other data aggregators, and measurement providers (such as Nielsen). This data provides information as to what content the consumer has watched. By understanding what a consumer has watched, it is possible to understand the consumer's taste and preferences, understand what TV shows they are watching, as well as when, where and on what device they are watching them. There are various deterministic methods (for example, understanding which member of the family has logged into their Netflix account) to determine which individual in a household is viewing which content.

With such a structure, the system is capable of integrating and processing data within different categories as well as across different categories. One example of this is comparing a complete set of user data for a given consumer to complete sets of data on other members of the market segment. Each complete user dataset is cross-compared to every other user dataset. The system then matches like behaviors, and is able to determine granular differences which may affect advertising performance. Such determinations then also fine tune the predictive algorithm on a consumer-by-consumer basis. In this way, the consumer data can be used in yield optimization, as further described herein, because it is possible to match audiences to TV campaigns.

## Setting Up an Advertising Campaign

An advertiser selects various campaign parameters and campaign goals. The advertiser can select particular parameters such as a demographic, and then values for a percentage of the overall population that meets the criteria of the campaign. For example, the advertiser can target women located in California between the ages of 20 and 30 years old, and specify a reach of 20% of that segment of the population. Another criterion could be the frequency by which the advertisement reaches a particular demographic, such as “two impression for each age group”. Criteria can be narrowed to identify users who are known to be in the market for a particular product, e.g., targeting women who have recently searched for a Nike® shoe.

Taken together, the advertiser-specified criteria can be ranked and weighted by importance. For example, women from San Francisco could be weighted more highly than women located in Sacramento; in which case, the system can allocate and budget impressions to those particular subsets of the overall demographic with higher weights. Additionally, special classifications can be placed on viewers based on their product purchase patterns and media consumption. Assumptions can be built into the campaign parameters based on purchase history, such that individuals who have purchased luxury handbags are more inclined to respond favorably to luxury handbag advertisements.

Next, the system makes it possible to run the various campaign parameters and goals against consumer data from first party databases, and received via third party API's.

The underlying consumer data inputs are integrated from a plurality of sources, which include internally developed datasets produced via a machine learning process, as well as third party data sources. For example, third party API's representing consumer and viewer data can be combined with data that exists as a result of internally processed data streaming from a series of processes that, for example, track viewer behavior and can predict viewer outcomes. Purchasing suggestions are provided, and comparisons can be derived based on the application of relevant, real-time metrics. The system can further assist in the execution of deal bidding and buying.

In a preferred embodiment, bidding and buying of advertising inventory is achieved at a much faster rate compared to existing methods due to the integration of data from a number of different sources within one system. This improvement in speed also provides significantly better accuracy in predicting data models around media consumption and consumer behavior. For example, the system is able to incorporate a range of data related to other buyers on the exchange, so that purchases are optimized based on considerations such as the distribution of inventory.

In a preferred embodiment, the system incorporates APIs from third parties who track relevant metrics such as consumer behavior and consumer demographics (for example, age, race, location, and gender). Relevant metrics are analyzed against the buyer's campaign requirements, such as budget, desired audience, and number of impressions.

In order to analyze available advertising inventory, the system intakes real-time inventory data via APIs from publishers and content providers. Data concerning inventory can be aggregated across mediums, so that inventory available for digital, mobile, TV, and OTT may be combined, thereby allowing advertisers to allocate their budget across a variety of mediums, device categories and content channels. Alternatively, such integration permits an MVPD to

allocate advertising content to platforms other than TV in the event that the content cannot be aligned with a favorable TV slot.

Once the advertiser has assigned campaign parameters, and the system has identified the possible inventories, the system permits the advertiser to choose a strategy for optimizing the allocation of impressions. There are various advertising strategies available to the advertiser.

Exemplary factors for advertising strategies are as follows:

Pacing: a rate at which an advertiser runs their advertising;

Uniform pacing: allocated evenly based on budget and length of campaign;

Accelerated buying: buying based on performance (e.g., the system detects a whole population of car video watchers and autonomously allocates based on that discovery)

Competitive pacing: if an advertiser's competitor is heavily buying a particular slot, the advertiser can choose whether to compete with them, or to allocate away from those slots (this is applicable to any device medium)

Specified time-frame strategies: advertising is purchased based on the time of day, and day of the week. End-user advertisers can buy inventories across the entire day, e.g., run every six hours, or limit the purchasers to only specified times during the day.

Inventory strategy: advertising is purchased based on maximizing performance of advertising dollars spent or meeting specified campaign parameters and goals.

Pricing strategy: purchasing focuses on staying within a defined budget. Budgets can be allocated by dollar amount according to inventory, medium, and/or time-frame.

Medium strategy: system detects which medium performs best to meet campaign goals.

Segmented media planning is the practice of deploying a number of media strategies to different inventory types. Existing strategies fold many consumers into high-level consumer classifications.

In one example of this process, an advertiser will set their media strategy directed towards a demographic defined as middle-aged women. The system can then segment the population of middle-aged women into various sub-segments, such as women with children, and women without children. Next, the system can match an advertising strategy, such as a uniform pacing strategy, to the individual level, say, a particular woman with children that has searched for diapers in the past hour. The system can then allocate the advertisement to a specific device of the woman, such as her mobile phone so that the next time she opens up an application, such as YouTube, the advert will display.

## The MVPD Environment

The technology described herein permits a MVPD to optimize an advertiser's campaign against available inventory, as depicted schematically in FIG. 8. The system is utilized by media owners who are sellers of a variety of advertising opportunities to marketers seeking to plan and deliver digital and TV advertising campaigns. The broadcaster, or MVPD, has a database 801 of customer data and has control over a system that matches advertising content to TV inventory. It accepts as input, one or more campaign descriptions 127 from the advertiser, the inventory of advertising slots 811, aggregated across a plurality of broadcast channels, and a set of clash constraints 821. Clash constraints are conditions on placement of advertising content that may come from advertisers themselves, regulatory

authorities, or from within the MVPD itself. These clash constraints are “hard” constraints in the sense that they cannot be violated by the final result of matching advertising content and a particular TV slot.

The MVPD, using the methods described herein, performs a match **831** of the campaigns across the slots in the inventory subject to the clash constraints. This produces a number of possible slots that might match a given campaign.

Those slots can then be subject to an interactive and/or iterative step of scenario modeling (either carried out automatically, or as a result of processing by a human operator) that utilizes information from the customer viewing and demographic data to obtain certain parameters, **0**. The matching process is subject to the “soft” clash constraints via, for example, a constraint satisfaction engine **851**.

The end process is one or more slots **861** that are deemed to be optimal for the campaign. At this stage, the campaign slots preferably include the best one, but it may have been necessary for tradeoffs to have been made.

#### Constraint Satisfaction

The MVPD environment can incorporate user-implied-and-weighted campaign constraints (e.g., hard constraints such as government regulations, and soft constraints relating to business goals) into the process of matching a campaign against inventory. The system produces revenue projections and campaign success rates by incorporating user-provided priority weights. The weight toggles and controls can be made available via a software interface, allowing a user to test the effects of higher or lower weights alongside a list of possible constraints. The system also produces an optimal weighting across soft constraints that maximizes the MVPD’s revenue from the sales and placement of advertisement in available advertising inventory slots across the various formats and screens controlled by the media owner. The advertising placement suggestions produced by the system effectively avoid scheduling conflicts by integrating and computing a variety of advertising placement constraints.

A hard constraint refers to legal (such as federal and state regulations) or otherwise mandatory limitations to placing advertising in particular time slots or within specified media. A mandatory limitation may be one imposed by an advertiser: for example, one manufacturer of soft drinks may require that its adverts do not run adjacently to adverts by specific competing manufacturers of soft drinks. A legal limitation may be a requirement that certain categories of advertisement do not run in conjunction with particular programming content. (For example, adverts for alcoholic beverages cannot be broadcast during childrens’ programming.)

A soft constraint refers to desired (non-mandatory) limitations to placing advertising in particular time slots within specified media.

Constraint satisfaction refers to the process of finding a solution to a set of constraints that impose conditions that the variables must satisfy. The solution therefore is a set of values for the variables that satisfies all constraints.

Algorithms which solve constraint satisfaction problems (“CSPs”) use some combination of search and inference (constraint propagation), where constraint propagation is used to reduce the search space, for example, in eliminating variable/value pairs which cannot belong to a solution. While advertising scheduling problems are often easily modeled by CSPs, this is only true for schedules that are only governed by hard constraints. When a schedule involves soft constraints (i.e., constraints that preferably should not be violated or that can be violated), CSPs do not

currently perform well. This is especially true when all solutions violate at least one soft constraint.

The technology herein makes it possible to solve advertising scheduling problems involving soft constraints by using one or more of partial CSPs (PCSPs) or weighted CSPs (WCSPs), in which an algorithm is used to minimize soft constraint violations. In general, CSPs do not easily incorporate secondary goals (i.e. a list of business goals ranked in order or preference). Where multiple goals with weighted preferences exist, the use of linear programming (LP) is generally more effective. Hence the underlying constraint satisfaction technology as used herein can use a combination of PCSPs and linear programming to integrate hard and soft constraints. The output is a suggestive placement of advertisements within available advertising inventory that maximizes revenue while satisfying a plurality of constraints. In this way it is possible to produce solutions that satisfy a subset of the total number of constraints.

While there is no general analytical solution for scheduling problems that have both hard constraints and soft constraints (or weights), an alternative way of formulating such problems is by a Weighted Constraint Satisfaction Problem (WCSP). Multiple approaches exist for solving WCSPs. For example, local stochastic search is effective for Max-SAT (see, e.g., the Internet site [www.cs.cornell.edu/selman/papers/pdf/maxsat.pdf](http://www.cs.cornell.edu/selman/papers/pdf/maxsat.pdf)). Evolutionary algorithms, or population based stochastic optimization (e.g., on the Internet at [ieeexplore.ieee.org/document/6900239/](http://ieeexplore.ieee.org/document/6900239/)) are also effective at solving a wider range of constraint satisfaction problems and can be used in conjunction with the methods herein. Conditional Preference Networks (CP-nets) are another technique for solving a mixture of hard and soft constraints (<http://www.math.unipd.it/~frossi/carmel-joh1.pdf>). WCSPs find more efficient solutions in several applications than LPs, for e.g. (see e.g., [www7.inra.fr/mia/T/degivry/Akplogan13.pdf](http://www7.inra.fr/mia/T/degivry/Akplogan13.pdf)).

As part of the underlying yield optimization solution, a sandbox environment (or equivalent environment that is shielded or isolated from external inputs) performs the following tasks, as shown in FIG. 8:

- a) Intakes data relating to advertising campaign details and advertising copy details, including any regulated broadcasting restriction codes;
- b) Intakes data variables and restrictions for the advertising inventories subject to the revenue optimization model;
- c) Produces an output to show the campaigns and advertising copies that meet the hard constraints; and
- d) Integrates, via a scenario modeler, user-selected weights to campaign completion, targeting, and priorities for a variety of campaigns.

The foregoing process allows the advertising inventory sellers to simulate the impact that different priority settings to constraints have on revenue returns. A user can manipulate toggles representing higher or lower weights on a variety of named constraints.

In one example, the underlying solution must place a major airline advertisement in a commercial break. To do this, the decision engine does not compute the value of an advertisement placement solely based on expected return and the fee earned per impression. Rather, the advertisement is processed using a linear combination of advertisement attributes and priority weights. The system computes all selections on rankings and places the airline’s advertisements based on the user-implied weights. This is an improvement over existing methods, which are only able to compute revenue based on fees earned per impression.



In a second example, the underlying solution has a hard constraint that prohibits the placement of a beauty product advertisement during a particular sitcom's commercial break. Placing this advertisement would violate the sitcom's programming restrictions, as delineated by the broadcaster of that program. To achieve a permitted placement and revenue-optimizing placement of the beauty product advertisement to avoid breaching the hard constraint, the system performs two stages of analysis. In the first stage, the engine filters out all advertising copy that violates unary constraints—one such unary constraint being program restrictions. In the second stage, the program produces the suggestive placement results for the beauty product advertisement via a front-end.

In a third example, an advertisement for a car manufacturer is placed in a commercial break. However, the content provider has a business policy that it does not air advertisements of competing companies in back to back advertising slots. This represents a soft constraint to the content provider. The system reviews the constraints provided by the content provider. The content provider-user sets a weight on the constraint, by using, for example, a toggle on the user interface that the user sets along a high-to-low slider bar. If the weight is set as a high priority, then the system will not produce any instance in which a competitor advertisement is placed alongside the car maker's advertisement. Further, the system can also integrate hard constraints at the same time. In this example, the user provides a hard constraint rule that prohibits alcohol and car advertisements from airing together. This hard constraint is simultaneously processed by the system to restrict ads promoting alcoholic beverages. The back-end system saves to memory restrictions on all subsequent incoming advertisement requests to ensure that both car manufacturers and alcoholic beverage advertisement requests are denied access to spots alongside the car maker's advertisements.

In all embodiments of the methods, the constraint satisfaction engine can solve for a plurality of constraint satisfaction problems. The solution resolves binary hard constraints in addition to non-binary soft constraints. In the third preferred embodiment involving the car maker's advertisement constraints, the engine's algorithm redirects the placement of non-authorized campaigns, for example, by placing advertisements by Guinness and Ford in slots other than the two framing the auto maker's advertisement. The engine can solve for any arbitrary set of hard and soft constraints.

In a fourth example, the engine considers which advertisements to play during the airing of a particular children's TV show. When determining which advertisement to place in a particular commercial break slot, the system first filters out all advertisements according the following implied constraints, and treats each constraint as a hard constraint:

1. Any advertisements that have reached their frequency cap: this would remove, for example, a car advertisement that has recently completed airing the full quantity of purchased airings.
2. Any advertisements that are forbidden under a regulation that place restrictions on airing soft drink advertisements during children's programming (such as those promulgated by the UK regulatory agency Clearcast).
3. Any advertisements that are forbidden under a regulation that places restrictions on airing alcoholic beverage advertisements during children's programming.

The above filters are programmatically applied by the system and the system produces suggestions on the highest yield advertisements available for placement.

In a fifth example, the system is able to interpret soft constraints that are advertising programming goals relating to the relative ordering of client advertisements. This occurs after hard constraint filters are applied and the data tables are no longer authorized to place non-permitted combinations. In this scenario, the system optimizes placements of advertisements according to business strategies relating to advertiser-client management. For example, the system places competitor advertisements as far apart as possible, without breaching any hard constraints. If the user-content provider wishes to favor one advertising client over another, they can adjust the soft constraint setting via a front end user interface to imply the preference. The system interprets the preference and places the more favored client's advertisements in programming breaks that optimize the client's campaign goals. Via the front-end, the user-content provider can rank advertiser-clients, and the system produces advertising scheduling suggestions according to the preference model. Optimizing Revenue from Virtual Sales of a Plurality of Advertising Inventory

An advantage of the technology herein is that a user can run a number of simulations of matching advertising campaigns with inventory slots. The simulations can thereby allow an optimization of yield in the sense of permitting the MVPD to assess what combination of matches will bring in the greatest amount of advertising revenue.

In one embodiment aggregated volumes of advertising inventory can be divided, such as programmatically, into commoditized inventory blocks that are then paired and targeted to the most relevant market segment for the respective inventory. The digital advertising yield is thereby optimized via segmentation and demand pairing in a manner that can be automated. The solution thereby optimizes the sales performance of inventory made available by content providers, such as Pay TV operators, TV OEMs, supply side platforms such as ad exchanges, ad networks, TV and media companies that sell advertising.

In another embodiment, an autonomous programmatic system enables media sellers and content providers to maximize total profits by selectively targeting inventory purchasers. The system includes a software interface via which sellers can access and manage a pool of available advertising inventory. The system may also include an API product that can be used to supply data into existing enterprise platforms on either or both sides of the transaction. The system can thereby be a stand-alone media management platform, or a data system that interacts with outside platforms.

While the system may be used by an MVPD to search for optimized placement of ads in TV inventory, the system can be used more generally in a variety of cross-screen settings. For that purpose, it intakes a list of available digital advertising inventories across a variety of device and media mediums. For example, in one embodiment, the system intakes a number, such as 1000, impressions on digital mobile applications, a number, such as 1000, of impressions from desktop websites, a number, such as 2, of TV spots on prime time programming, and a further number, such as 10, spots of radio content. The system aggregates all possible sales-side inventory into one common relational database. The data can be structured and tagged according to its relevant attributes, such as media medium, inventory time schedule, and media broadcast. The structuring and tagging can be done automatically based on certain rules and ability to recognize the content.

The system, which can be the same system, integrates data related to the current prevailing market conditions such as current clearance prices, winning bid prices, and the

number of demand partners (purchasers of the inventory). Additional market data inputs may be integrated into the database. These can include historical characteristics of each market to determine appropriate pricing models.

Via a separate demand-side data process, the system collects and aggregates demand side (inventory buyer) data. Demand side data can originate through demand side partners via their APIs. Demand side data lists all current and active requests for available advertising inventory. The system then runs an optimization process on the demand-side data, dividing the inventory into smaller saleable blocks of data. The blocks of data are ranked and rated to show estimates on how much each block will produce in revenue to the supply-side sellers. The optimization process can also map the target segments from the demand side to available advertising inventory in the supply-side dataset. In one example, a demand-side target segment shows that a purchaser of inventory is interested in buying 1000 impressions from 29 year-old men. The system identifies the top value per impression inventory from the supply-side dataset and automatically pushes the demand data to the supply-side user.

Once paired, the system provides suggestions to supply-side users via an interface. These supply-side users include content providers, OEMs, supply side platforms such as ad exchanges, ad networks, TV and media companies that sell advertising such as Fox News Agency, CNN, SKY TV. Alternatively, the system may allocate through an API various blocks of inventory to be sold in a variety of inventory markets. The API's data is optimized using the methods noted above, and updated automatically as new inventories become available. The solution's data output is utilized to maximize total revenue as well as provide the best match to target segments.

As new inventory markets become known and identified, the optimization tool can take into account relevant characteristics of these markets in the solution optimization process. As new media conduits become available, such as the creation of new media consumption devices, the system will optimize around information regarding consumer use of these devices. As new demand partners are added, the system will re-calculate the allocation of inventory sales based on the aggregate demand-side marketplace.

Unlike current systems, which are not able to autonomously account for individual nuances within demand markets, or automatically segment the demand markets into granular buckets of analysis, the current system benefits from the automated tagging, and continuous analysis and re-organization of available data inputs. The system processes a large number of possible monetization opportunities (including preferably exhausting them all), and proposes estimates on which opportunities are the most economically rewarding.

#### Delivering and Optimizing Cross-Screen Advertising Content

The technology described herein also permits an advertiser to target advertising content to a consumer across more than one media conduit, including both TV and online media. This is particularly relevant in a situation where an MVPD is unable to find an optimum placement for advertising content in available TV slots but has access to inventory on other devices known to be used by the relevant consumer demographic.

There are two types of environment in which an advertiser can target a consumer. In a 1:1 environment, a DSP can just use the actual segment and/or a modeled out version of the actual segment, to make a real time decision to place the

advert if the consumer matches the targeting parameters. In an index approach, when it is not possible to target 1:1 and it is not possible to do dynamic advert insertion or real time decisioning, the system instead looks at concentration of viewers projected to access the slot (such as a TV program or VOD program) and then targets the slots that have the highest concentration of the target consumers.

In a preferred embodiment, the advertiser has control over the allocation of the advertising content because the advertiser accesses the system via a unified interface that presents information about inventory, manages bids on the inventory, and provides a list of potential advertising targets consistent with a campaign description and the advertiser's budget. The system then communicates with, for example, supply-side providers to ensure that the desired slots are purchased, typically via a bidding process, and the advertising content is delivered or caused to be delivered.

In one embodiment, the technology provides for an advertising campaign that involves delivery of content across two or more media conduits, rather than delivery of a single advertisement to multiple consumers at different times on, say, TV only. The system thereby permits delivery of advertising content to a given consumer, or a population of consumers, on more than one device. For example, a consumer may view a portion of the campaign on a TV, and may also see the campaign within a desktop browser session on their laptop or on their OTT device. In this context, the TV inventory can be purchased across a variety of TV consumption devices that include, but are not limited to linear, time-shifted, traditional and programmatic TV, according to bid methodology described herein or familiar to those skilled in the art. In some instances, the advertiser desires to cap the number of impressions that a given consumer receives; in other instances, the advertiser wants to extend the campaign from one media to another based on metrics calculated across various media conduits. The method permits the advertiser to target segments of a population more precisely than before, as well as achieve fine scale refinement of a campaign based on performance indicators from more than one conduit.

There are two aspects of the technology that enable an advertiser to successfully manage and refine an advertising campaign: the system is able to keep track of which devices a given consumer can access, as well as on which devices the user has already been exposed to the advertising campaign; the system can also identify those consumers that are most likely to be interested in the campaign's content. Accuracy in targeting can thus be achieved via predictions based on a mapping of consumer behavior from aggregated cross-screen viewership data.

The analytics portion of the system is capable of accepting unlimited data inputs regarding consumer behavior across various media. The system uses this data to optimize consumer classifications. A second part of the output is improved measurement and prediction of future consumer behaviors based on the data on cross-screen behavior.

Analysis of cross-screen data is able to determine where and when a consumer has viewed an advertisement, or a particular version of it, and thereby permits advertisers to schedule broadcast of an advertising campaign across multiple platforms. Advertisers can then schedule where, when, and how an advertisement is subsequently broadcast. They have control over retargeting (whether they show the same advertising more than once), or can choose to broadcast a multi-chapter advertising story.

One method for managing delivery of advertising content to a consumer across two more display devices is illustrated

in FIG. 7. A consumer graph is, or has been, constructed **710**, or is continually under construction and revision, according to methods described elsewhere herein, and a pool of consumers is defined **730**, based on the graph of consumer properties, wherein the graph contains information about the devices used by each consumer and demographic data on each consumer, and wherein the pool of consumers contains consumers having at least a threshold similarity to a member of a target audience.

The system receives a list of advertising inventory **712** from one or more media conduits or content providers, wherein the list of inventory comprises one or more slots for TV and online.

The system receives a pricepoint **702** one or more campaign descriptions **705** from an advertiser, wherein each of the campaign descriptions **705** comprises a schedule for delivery of or more items of advertising content across two or more devices accessed by a consumer, and a target audience **720**, wherein the target audience is defined by one or more demographic factors selected from: age range, gender, and location. Pricepoint **702** represents an advertiser's budget for the advertising campaign. The budget can be allocated across multiple slots, and across multiple media conduits, according to the inventory and goals for the campaign. Goals may include the target audience desired to be reached, and the hoped for number of impressions.

Based on the pool of consumers, the campaign descriptions and available inventory, the system can identify one or more advertising targets, wherein each of the one or more advertising targets comprises two or more slots, consistent with a given pricepoint **702** associated with a campaign description **705**. It is then possible to allocate the advertising content of the one or more campaign descriptions to the one or more advertising targets based on the inventory.

The foregoing steps may be carried out in orders other than as described above, or iteratively, sequentially, or simultaneously, in part. Thus, the system may receive the campaign descriptions and pricepoint(s), at the same time as it receives advertising inventory, or beforehand, or afterwards. The consumer graph, additionally, may be continually being updated.

For a given consumer, a number of devices accessed by that consumer are identified **740**. This can be from constructing a device graph as discussed elsewhere herein. Those consumers for which more than one device has been associated, can be targeted by an advertising campaign herein.

Inputs to the system of the various categories of data (inventory, advertising campaigns, etc.) can be via various application program interfaces (APIs), the development of which is within the capability of one skilled in the art.

Then **770**, for each slot in an advertising target, the system makes a bid on the slot consistent with the pricepoint; for two slots where a bid is a winning bid, the system then instructs a first content provider to deliver a first item of advertising content in a first slot to a pool of consumers on a first device, and for a second slot, can then instruct a second content provider to deliver a second item of advertising content in the second slot on a second device. It is preferable that at least one of the first device and the second device is a TV.

It is to be understood that the instructing and delivery steps are optional for a given entity performing the method, once the slots of TV and online inventory have been identified as consistent with the advertising campaign.

The methods herein can also be used to optimize an advertising campaign across a plurality of devices accessible

to a consumer. Such methods build upon methods of delivery described hereinabove and with respect to FIG. 7. Once it has been determined that a consumer is a member of a target audience, and a first and second device accessible to the consumer have been identified, an advertiser wants to purchase of slots for a first and second item of advertising content on the first and second devices, consistent with an advertising budget and the target audience, in a manner that improves upon prior campaigns.

In this instance, the system can receive feedback on a consumer's response to the first and second items of advertising content, and based on that information as well as similar information from other consumers, it is possible to use the feedback to instruct purchase of further slots for the first and second items of advertising content.

For example, the system can receive a first datum from a first tag that accompanied the first item of advertising content to validate whether a particular consumer viewed the first item of advertising content on the first device as well as a second datum from a second tag that accompanied the second item of advertising content. A given datum of content can be a beacon, such as communicated via a protocol such as VPAID or VAST.

In some embodiments, a datum can comprise a confirmation of whether the consumer has seen the first item of advertising content, in which case the second item of advertising content is not delivered to the consumer until the consumer has seen the first item of advertising content.

In some embodiments, the advertising campaign can be optimized in a number of different ways. Although measurements of deduplicated reach, as described elsewhere herein, can be used to assess—and refine—the effectiveness of an advertising campaign, another factor is the overall cost effectiveness of the campaign. For example, given a budget, or a dollar-amount spent per advertisement, the cost per impression can be calculated. This number can be optimized over successive iterations of the campaign.

In other embodiments, an advertising campaign is updated and optimized during its own term. For example, a campaign may be scheduled to run over a particular time period, such as 3 days, 1 week, 2 weeks, 1 month, or 3 months. The system herein can provide feedback on the efficacy of the campaign before it is complete, and can therefore provide an advertiser with an ability and opportunity to adjust parameters of the campaign in order to improve its reach. Such parameters include, but are not limited to, aspects of audience demographic such as age, income, location, and media on which the advertisement is delivered, such as TV station, or time of day.

The system and methods herein can still further provide an advertiser with a way to project viewer data accumulated historically on to future potential viewing habits, for example, using look-alike modelling. The historical data can include data acquired during the course of the campaign.

#### 55 Computational Implementation

The computer functions for manipulations of advertising campaign data, advertising inventory, and consumer and device graphs, in representations such as bit-strings, as well as optimizing yield and working with hard and soft constraints, can be developed and implemented by a programmer or a team of programmers skilled in the art. The functions can be implemented in a number and variety of programming languages, including, in some cases mixed implementations. For example, the functions as well as scripting functions can be programmed in functional programming languages such as: Scala, Golang, and R. Other programming languages may be used for portions of the

implementation, such as Prolog, Pascal, C, C++, Java, Python, VisualBasic, Perl, .Net languages such as C#, and other equivalent languages not listed herein. The capability of the technology is not limited by or dependent on the underlying programming language used for implementation or control of access to the basic functions. Alternatively, the functionality could be implemented from higher level functions such as tool-kits that rely on previously developed functions for manipulating mathematical expressions such as bit-strings and sparse vectors.

The technology herein can be developed to run with any of the well-known computer operating systems in use today, as well as others, not listed herein. Those operating systems include, but are not limited to: Windows (including variants such as Windows XP, Windows95, Windows2000, Windows Vista, Windows 7, and Windows 8, Windows Mobile, and Windows 10, and intermediate updates thereof, available from Microsoft Corporation); Apple iOS (including variants such as iOS3, iOS4, and iOS5, iOS6, iOS7, iOS8, and iOS9, and intervening updates to the same); Apple Mac operating systems such as OS9, OS 10.x (including variants known as "Leopard", "Snow Leopard", "Mountain Lion", and "Lion"; Android operating systems; the UNIX operating system (e.g., Berkeley Standard version); and the Linux operating system (e.g., available from numerous distributors of free or "open source" software).

To the extent that a given implementation relies on other software components, already implemented, such as functions for manipulating sparse vectors, and functions for calculating similarity metrics of vectors, those functions can be assumed to be accessible to a programmer of skill in the art.

Furthermore, it is to be understood that the executable instructions that cause a suitably-programmed computer to execute the methods described herein, can be stored and delivered in any suitable computer-readable format. This can include, but is not limited to, a portable readable drive, such as a large capacity "hard-drive", or a "pen-drive", such as connects to a computer's USB port, an internal drive to a computer, and a CD-Rom or an optical disk. It is further to be understood that while the executable instructions can be stored on a portable computer-readable medium and delivered in such tangible form to a purchaser or user, the executable instructions can also be downloaded from a remote location to the user's computer, such as via an Internet connection which itself may rely in part on a wireless technology such as WiFi. Such an aspect of the technology does not imply that the executable instructions take the form of a signal or other non-tangible embodiment. The executable instructions may also be executed as part of a "virtual machine" implementation.

The technology herein is not limited to a particular web browser version or type; it can be envisaged that the technology can be practiced with one or more of: Safari, Internet Explorer, Edge, FireFox, Chrome, or Opera, and any version thereof.

#### Computing Apparatus

An exemplary general-purpose computing apparatus **900** suitable for practicing the methods described herein is depicted schematically in FIG. 9. Such a computer apparatus can be located within the control of an MVPD, such as be linked to an intranet within an MVPD corporate environment.

The computer system **900** comprises at least one data processing unit (CPU) **922**, a memory **938**, which will typically include both high speed random access memory as well as non-volatile memory (such as one or more magnetic

disk drives), a user interface **924**, one more disks **934**, and at least one network or other communication interface connection **936** for communicating with other computers over a network, including the Internet, as well as other devices, such as via a high speed networking cable, or a wireless connection. There may optionally be a firewall **952** between the computer and the Internet. At least the CPU **922**, memory **938**, user interface **924**, disk **934** and network interface **936**, communicate with one another via at least one communication bus **933**.

CPU **922** may optionally include a vector processor, optimized for manipulating large vectors of data.

Memory **938** stores procedures and data, typically including some or all of: an operating system **940** for providing basic system services; one or more application programs, such as a parser routine **950**, and a compiler (not shown in FIG. 9), a file system **942**, one or more databases **944** that store advertising inventory **946**, campaign descriptions **948**, and other information, and optionally a floating point coprocessor where necessary for carrying out high level mathematical operations. The methods of the present invention may also draw upon functions contained in one or more dynamically linked libraries, not shown in FIG. 9, but stored either in memory **938**, or on disk **934**.

The database and other routines shown in FIG. 9 as stored in memory **938** may instead, optionally, be stored on disk **934** where the amount of data in the database is too great to be efficiently stored in memory **938**. The database may also instead, or in part, be stored on one or more remote computers that communicate with computer system **900** through network interface **936**.

Memory **938** is encoded with instructions for receiving input from one or more advertisers and for calculating a similarity score for consumers against one another. Instructions further include programmed instructions for performing one or more of parsing, calculating a metric, and various statistical analyses. In some embodiments, the sparse vector themselves are not calculated on the computer **900** but are performed on a different computer and, e.g., transferred via network interface **936** to computer **900**.

Various implementations of the technology herein can be contemplated, particularly as performed on computing apparatuses of varying complexity, including, without limitation, workstations, PC's, laptops, notebooks, tablets, netbooks, and other mobile computing devices, including cell-phones, mobile phones, wearable devices, and personal digital assistants. The computing devices can have suitably configured processors, including, without limitation, graphics processors, vector processors, and math coprocessors, for running software that carries out the methods herein. In addition, certain computing functions are typically distributed across more than one computer so that, for example, one computer accepts input and instructions, and a second or additional computers receive the instructions via a network connection and carry out the processing at a remote location, and optionally communicate results or output back to the first computer.

Control of the computing apparatuses can be via a user interface **924**, which may comprise a display, mouse **926**, keyboard **930**, and/or other items not shown in FIG. 9, such as a track-pad, track-ball, touch-screen, stylus, speech-recognition, gesture-recognition technology, or other input such as based on a user's eye-movement, or any subcombination or combination of inputs thereof. Additionally, implementations are configured that permit a purchaser of advertising inventory to access computer **900** remotely, over a network

connection, and to view inventory via an interface having attributes comparable to interface 924.

In one embodiment, the computing apparatus can be configured to restrict user access, such as by scanning a QR-code, gesture recognition, biometric data input, or password input.

The manner of operation of the technology, when reduced to an embodiment as one or more software modules, functions, or subroutines, can be in a batch-mode—as on a stored database of inventory and consumer data, processed in batches, or by interaction with a user who inputs specific instructions for a single advertising campaign.

The results of matching advertising inventory to criteria for an advertising campaign, as created by the technology herein, can be displayed in tangible form, such as on one or more computer displays, such as a monitor, laptop display, or the screen of a tablet, notebook, netbook, or cellular phone. The results can further be printed to paper form, stored as electronic files in a format for saving on a computer-readable medium or for transferring or sharing between computers, or projected onto a screen of an auditorium such as during a presentation.

Toolkit: The technology herein can be implemented in a manner that gives a user (such as a purchaser of advertising inventory) access to, and control over, basic functions that provide key elements of advertising campaign management. Certain default settings can be built in to a computer-implementation, but the user can be given as much choice as possible over the features that are used in assigning inventory, thereby permitting a user to remove certain features from consideration or adjust their weightings, as applicable.

The toolkit can be operated via scripting tools, as well as or instead of a graphical user interface that offers touch-screen selection, and/or menu pull-downs, as applicable to the sophistication of the user. The manner of access to the underlying tools by a user is not in any way a limitation on the technology's novelty, inventiveness, or utility.

Accordingly, the methods herein may be implemented on or across one or more computing apparatuses having processors configured to execute the methods, and encoded as executable instructions in computer readable media.

For example, the technology herein includes computer readable media encoded with instructions for executing a method for allocating delivery of video advertising content to a consumer on a TV, the instructions including: instructions for receiving a pricepoint and one or more campaign descriptions from an advertiser, wherein each of the campaign descriptions comprises a schedule for delivery of an item of advertising content across one or more TVs accessed by a consumer, and a target audience, wherein the target audience is defined by one or more demographic factors; instructions for identifying one or more hard constraints associated with the one or more campaign descriptions; instructions for defining a pool of consumers based on a graph of consumer properties, wherein the graph contains information about two or more TV and mobile devices used by each consumer, demographic and online behavioral data on each consumer and similarities between pairs of consumers, and wherein the pool of consumers comprises consumers having at least a threshold similarity to a member of the target audience; instructions for receiving a list of inventory from one or more content providers, wherein the list of inventory comprises one or more slots for TV and online; instructions for identifying one or more advertising targets, wherein each of the one or more advertising targets comprises a sequence of slots consistent with one or more of the campaign descriptions and the one or more hard constraints,

and has an overall cost consistent with the pricepoint; instructions for carrying out an optimization of allocation of the advertising content of the one or more campaign descriptions to the one or more advertising targets based on one or more soft constraints, thereby producing one or more solutions; instructions for communicating a list of one or more solutions to the advertiser, wherein a solution comprises a match of the campaign description to one or more slots within TV content identified as likely to be viewed by the pool of consumers; and instructions for delivering the item of advertising content to a consumer in the pool of consumers via a first media conduit on a television.

Correspondingly, the technology herein also includes a computing apparatus having at least one processor configured to execute instructions for implementing a method for allocating delivery of video advertising content to a consumer on a TV, the instructions including: instructions for receiving a pricepoint and one or more campaign descriptions from an advertiser, wherein each of the campaign descriptions comprises a schedule for delivery of an item of advertising content across one or more TVs accessed by a consumer, and a target audience, wherein the target audience is defined by one or more demographic factors; instructions for identifying one or more hard constraints associated with the one or more campaign descriptions; instructions for defining a pool of consumers based on a graph of consumer properties, wherein the graph contains information about two or more TV and mobile devices used by each consumer, demographic and online behavioral data on each consumer and similarities between pairs of consumers, and wherein the pool of consumers comprises consumers having at least a threshold similarity to a member of the target audience; instructions for receiving a list of inventory from one or more content providers, wherein the list of inventory comprises one or more slots for TV and online; instructions for identifying one or more advertising targets, wherein each of the one or more advertising targets comprises a sequence of slots consistent with one or more of the campaign descriptions and the one or more hard constraints, and has an overall cost consistent with the pricepoint; instructions for carrying out an optimization of allocation of the advertising content of the one or more campaign descriptions to the one or more advertising targets based on one or more soft constraints, thereby producing one or more solutions; instructions for communicating a list of one or more solutions to the advertiser, wherein a solution comprises a match of the campaign description to one or more slots within TV content identified as likely to be viewed by the pool of consumers; and instructions for delivering the item of advertising content to a consumer in the pool of consumers via a first media conduit on a television.

#### Cloud Computing

The methods herein can be implemented to run in the "cloud." Thus the processes that one or more computer processors execute to carry out the computer-based methods herein do not need to be carried out by a single computing machine or apparatus. Processes and calculations can be distributed amongst multiple processors in one or more datacenters that are physically situated in different locations from one another. Data is exchanged with the various processors using network connections such as the Internet. Preferably, security protocols such as encryption are utilized to minimize the possibility that consumer data can be compromised. Calculations that are performed across one or more locations remote from an entity such as a MVPD or a DSP include calculation of consumer and device graphs, and updates to the same.

## Example 1: User Interface

An exemplary user interface is shown in FIGS. 10A-10D, in which successive panels illustrate an interface for successive steps in the workflow of yield optimization, such as may be practiced within an MVPD environment.

In FIG. 10A, campaign details are uploaded to the system via a front-end.

In FIG. 10B, a user-front end is provided to organize hard constraints associated with available advertising inventory spots.

In FIG. 10C, a user-front end is provided to display advertising spots that do not violate any of the hard constraints in FIG. 10B.

In FIG. 10D, a front-end displays the system's scenario modeling interface with weighted implied soft constraints and predicted revenue under selected weights.

All references cited herein are incorporated by reference in their entireties.

The foregoing description is intended to illustrate various aspects of the instant technology. It is not intended that the examples presented herein limit the scope of the appended claims. The invention now being fully described, it will be apparent to one of ordinary skill in the art that many changes and modifications can be made thereto without departing from the spirit or scope of the appended claims.

What is claimed:

1. A method, comprising:

receiving first viewership data associated with a first consumer view of an advertisement of an advertising campaign, wherein the first consumer view of the advertisement includes a first view of a first portion of the advertisement on a first device and a second view of a second portion of the advertisement on a second device;

generating a first consumer classification based on the first viewership data;

calculating a predictive model of consumer behavior in view of the first consumer classification and a plurality of soft constraints, including running a first plurality of simulations using the predictive model in a plurality of advertisement campaigns, wherein running the first plurality of simulations includes determining a constraint satisfaction solution relative to a particular advertising campaign, and adjusting at least one soft constraint of the plurality of soft constraints in response to a user input;

receiving second viewership data associated with a second consumer view of the advertisement of the advertising campaign;

generating a second consumer classification based on the second viewership data;

updating the predictive model in view of the second consumer classification and in view of running a second plurality of simulations using the second consumer classification or second viewership data;

calculating an accuracy of the predictive model;

calculating lift data of the advertising campaign;

comparing the lift data to the predictive model;

updating the predictive model in view of the compared lift data and the calculated accuracy of the predictive model; and

sending the results of at least one of the first consumer classification, the second consumer classification, the predictive model, or the lift data, to a user interface.

2. The method of claim 1, wherein at least the first consumer classification is a multi-dimensional consumer classification, and where at least one dimension of the first consumer classification is associated with a device on which the advertisement was viewed.

3. The method of claim 2, wherein the multi-dimensional consumer classification includes linked data from each dimension of the first consumer classification, wherein each dimension represents a different device the first consumer used to view the advertisement.

4. The method of claim 2, wherein the device of the first consumer includes at least one of a TV, a personal computer, or a mobile device.

5. The method of claim 1, wherein the first consumer and the second consumer are a subset of a total audience of the advertising campaign.

6. The method of claim 1, wherein the predictive model is updated by calculating a difference between the predictive model of consumer behavior and actual consumer behavior, using a feedback loop including an input of at least the first viewership data, the second viewership data, or a third viewership data, and generating an output of at least an updated predictive model of consumer behavior.

7. The method of claim 6, wherein the predictive model is updated in real-time.

8. The method of claim 1, further comprising deduplicating the first viewership data of the first consumer in response to the first consumer viewing the first portion of the advertisement on the first device and viewing the second portion of the advertisement on the second device.

9. The method of claim 8, wherein the first portion of the advertisement includes a segment of the advertisement that comprises less than a full length of a runtime of the advertisement or a subset of advertisements of a plurality of advertisements intended to be viewed by the first consumer in the advertising campaign.

10. The method of claim 1, wherein the lift data includes at least one of brand lift data or sales lift data.

11. A system, comprising:

at least one computer-readable media configured to store instructions; and

at least one processor coupled to the computer-readable media, the at least one processor configured to execute the instruction to cause the processor to perform operations, the operations comprising:

receive first viewership data associated with a first consumer view of an advertisement of an advertising campaign, wherein the first consumer view of the advertisement includes a first view of a first portion of the advertisement on a first device and a second view of a second portion of the advertisement on a second device;

generate a first multi-dimensional consumer classification based on the first viewership data, the first multi-dimensional consumer classification including at least one dimension associated with a device on which the advertisement was viewed;

calculate a predictive model of consumer behavior in view of the first consumer classification, the first consumer classification being within a particular audience segment;

receive second viewership data associated with a second consumer view of the advertisement of the advertising campaign;

41

generate a second consumer classification based on the second viewership data, the second consumer classification being within the particular audience segment;

update the predictive model in view of the second consumer classification;

determine, in view of a plurality of soft constraints, a particular configuration of the plurality of soft constraints to satisfy a constraint satisfaction relative to the advertisement campaign, the particular solution including an adjustment to at least on soft constraint of the plurality of soft constraints in response to a user input;

calculate an accuracy of the predictive model in view of the particular configuration of the plurality of soft constraints;

calculate lift data of the advertising campaign;

compare the lift data to the predictive model;

update the predictive model in view of the compared lift data and the calculated accuracy of the predictive model; and

send the results of at least one of the first consumer classification, the second consumer classification, the predictive model, the particular audience segment, or the lift data, to a user interface.

**12.** The system of claim **11**, wherein the multi-dimensional consumer classification includes linked data from each dimension of the first consumer classification, wherein each dimension represents a different device the first consumer used to view the advertisement.

**13.** The system of claim **11**, wherein the device of the first consumer includes at least one of a TV, a personal computer, or a mobile device.

**14.** The system of claim **11**, wherein the first consumer and the second consumer are a subset of a total audience of the advertising campaign.

**15.** The system of claim **11**, wherein the predictive model is updated by calculating a difference between the predictive model of consumer behavior and actual consumer behavior, using a feedback loop including an input of at least the first viewership data, the second viewership data, or a third viewership data, and generating an output of at least an updated predictive model of consumer behavior.

**16.** The system of claim **15**, wherein the predictive model is updated in real-time.

**17.** The system of claim **11**, further comprising deduplicate the first viewership data of the first consumer in response to the first consumer viewing the first portion of the advertisement on first device and viewing the second portion of the advertisement on the second device.

**18.** The system of claim **17**, wherein the first portion of the advertisement includes a segment of the advertisement that comprises less than a full length of a runtime of the advertisement or a subset of advertisements of a plurality of advertisements intended to be viewed by the first consumer in the advertising campaign.

42

**19.** The system of claim **11**, wherein the lift data includes at least one of brand lift data or sales lift data.

**20.** A non-transitory computer readable medium having stored therein executable code that, when executed by a processor, cause the processor to perform operations comprising:

receiving first viewership data associated with a first consumer view of an advertisement of an advertising campaign, the advertising campaign having a plurality of soft constraints, wherein the first consumer view of the advertisement includes a first view of a first portion of the advertisement on a first device and a second view of a second portion of the advertisement on a second device;

deduplicating the first viewership data of the first consumer in response to the first consumer viewing a first portion of the advertisement on a first device and viewing a second portion of the advertisement on a second device;

generating a first consumer classification based on the deduplicated first viewership data;

adjusting, in response to a user input, at least one of the plurality of soft constraints to satisfy a constraint satisfaction relative to the advertising campaign;

calculating a predictive model of consumer behavior in view of the first consumer classification and in view of an adjustment to the at least one of the plurality of soft constraints;

receiving second viewership data associated with a second consumer view of the advertisement of the advertising campaign;

deduplicating the second viewership data of the second consumer in response to the second consumer viewing the first portion of the advertisement on a third device and viewing a second portion of the advertisement on a fourth device;

generating a second consumer classification based on the deduplicated second viewership data;

updating the predictive model in view of the second consumer classification;

calculating an accuracy of the predictive model;

calculating lift data of the advertising campaign;

comparing the lift data to the predictive model;

updating the predictive model in view of the compared lift data and the calculated accuracy of the predictive model;

generating an incremental reach after running the advertising campaign; and

sending the results of at least one of the first consumer classification, the second consumer classification, the predictive model, the incremental reach, or the lift data, to a user interface.

\* \* \* \* \*